

Adaptive empirical Bayesian smoothing splines

Paulo Serra and Tatyana Krivobokova

Georg-August-Universität Göttingen

November 18, 2015

Abstract

In this paper we develop and study adaptive empirical Bayesian smoothing splines. These are smoothing splines with both smoothing parameter and penalty order determined via the empirical Bayes method from the marginal likelihood of the model. The selected order and smoothing parameter are used to construct adaptive credible sets with good frequentist coverage for the underlying regression function. We use these credible sets as a proxy to show the superior performance of adaptive empirical Bayesian smoothing splines compared to frequentist smoothing splines.

Keywords: Adaptive estimation; Unbiased risk minimiser; Maximum likelihood; Oracle parameters

1 Introduction

Consider n observations from the non-parametric regression model

$$Y_i = f(x_i) + \sigma\epsilon_i, \quad i = 1, \dots, n. \quad (1)$$

The function f is assumed to belong to a Sobolev class $\mathcal{W}_\beta(M)$, a collection of continuous functions $f \in L_2$ such that $f^{(\beta-1)}$ is absolutely continuous and $\|f^{(\beta)}\|^2 < M^2$, where $\|\cdot\|$ is the ℓ_2 -norm. The design points $\mathbf{x} = (x_1, \dots, x_n) \in [0, 1]^n$ are $x_i = (2i-1)/(2n)$, the observation errors $\epsilon_1, \dots, \epsilon_n$ are assumed to be i.i.d. standard Gaussian random variables and $\sigma^2 > 0$. Parameters f , β , and σ^2 are unknown and of interest.

In this paper we study a smoothing spline estimator for f , which is the unique minimiser in \mathcal{W}_q of the penalised least squares criterion

$$\frac{1}{n} \sum_{i=1}^n \{Y_i - f(x_i)\}^2 + \lambda \int_0^1 \{f^{(q)}(t)\}^2 dt, \quad \lambda > 0, \quad q \in \mathbb{N} \quad (2)$$

and is well known to be a natural polynomial smoothing spline of degree $2q-1$ with knots at the observation points; see Wahba (1990).

The performance of smoothing splines as data-smoothers crucially depends on the choice of the smoothing parameter λ , which balances fidelity to the data and smoothness

of the estimator. Criteria to select such a smoothing parameter can be obtained under two paradigms, which correspond to making different assumptions on the data-generating mechanism. One possibility is to assume that the regression function f is some fixed function from a certain class (frequentist model). In this case λ is estimated by minimising an unbiased estimator of the mean integrated squared error (unbiased risk estimator). Generalised cross validation (GCV), Mallows's C_p and Akaike's information criterion are all asymptotically equivalent criteria of this type. In the following $\hat{\lambda}_f$ denotes a minimiser of one of these criteria.

Another possibility is to assume that the regression function f is a realisation of some stochastic process (Bayesian model). Here a conjugate prior is put on the regression function, such that the resulting posterior mean coincides with the smoothing spline estimator. The smoothing parameter λ is then a so-called *hyper-parameter* of the prior. Its estimator is set to a maximiser of the resulting marginal likelihood (empirical Bayes method) and will be denoted by $\hat{\lambda}_q$.

The prior we use is a Gaussian prior. For these, conjugacy properties are often explored to directly study the posterior and specific Bayes estimators. Properties of Gaussian process priors (not necessarily conjugate) can be found in van der Vaart and van Zanten (2008); modifications to obtain adaptive priors were proposed in van der Vaart and van Zanten (2009). A (by no means extensive) list of results on adaptation using Gaussian priors in regression and the closely related Gaussian white noise model include Belitser and Ghosal (2003); Babenko and Belitser (2010); Knapik et al. (2011, 2013); de Jonge and van Zanten (2010, 2012); Szabó et al. (2014).

Bayesian smoothing splines with Gaussian priors have been first considered in Kimeldorf and Wahba (1970). Extensions and modifications of these splines have been discussed e.g. in Kohn and Ansley (1987), Speckman and Sun (2003) or Yue et al. (2014).

The asymptotic distributions of $\hat{\lambda}_f$ and $\hat{\lambda}_q$ can be computed under the assumption that the data come from the frequentist model with f as a fixed, “true” regression function of interest. This allows a direct comparison of these two estimators obtained under different paradigms. Krivobokova (2013) shows that $\hat{\lambda}_f$ and $\hat{\lambda}_q$ are consistent for certain oracles and that the asymptotic variance of $\hat{\lambda}_f$ can be several times larger than that of $\hat{\lambda}_q$.

The literature on adaptive Bayesian nonparametric estimation in non-parametric regression, and their frequentist performance is already quite extensive. For general priors, sufficient conditions for so called posterior contraction were proposed in Ghosal et al. (2000); Shen and Wasserman (2001); Ghosal and van der Vaart (2007) – posterior contraction at a given rate ensures the existence of frequentist estimators with the same rate. Adaptation is usually achieved by considering a family of priors indexed by a hyper-parameter (like λ and q above). If the regression function f belongs to a given smoothness class and there is some value of the hyper-parameter such that the resulting posterior contracts at the minimax rate for f in that class, then either endowing the hyper-parameter with a prior (hierarchical Bayes), or picking the hyper-parameter in a data-driven way (empirical Bayes), can lead to posteriors that contract adaptively. This approach has been used in Belitser and Levit (2003); Zhang (2005); Johnstone and Silverman (2005); McAuliffe et al. (2006); Ghosal et al. (2008); Knapik et al. (2013); Shen and Ghosal (2015), among

others. Here we consider the empirical Bayes approach.

Smoothing parameters that minimise an unbiased risk estimator (e.g., GCV) are predominant in practice and known to have good theoretical properties. In particular, if there is a mismatch between the order of the spline ($2q - 1$) and the smoothness of the regression function (f admits more than just q square integrable derivatives), then $\hat{\lambda}_f$ adapts to this extra smoothness, but only up to $2q$. In contrast, $\hat{\lambda}_q$ does not adapt and its rate is determined by q only. The main question that we address in this paper is whether one can obtain a *selector* \hat{q} , such that the resulting $\hat{\lambda}_{\hat{q}}$ not only adapts to the underlying smoothness of f , but also outperforms $\hat{\lambda}_f$ due to a much smaller variance.

To derive such a selection criterion for q we use the fact that the prior distribution depends on q , albeit in an implicit way, and apply the empirical Bayes approach. Contrary to the selection of the smoothing parameter λ , the selection of the order of smoothing splines has received very little attention in the literature.

Since $\hat{\lambda}_f$ and $\hat{\lambda}_{\hat{q}}$ are associated with splines of different order, direct comparison between the two smoothing parameters is not adequate. Instead, we construct credible ℓ_2 balls with good frequentist coverage, obtained from a high probability region of the posterior corresponding to hyper-parameters \hat{q} and $\hat{\lambda}_{\hat{q}}$ selected via empirical Bayes. Subsequently, we show that if the center of this ball is replaced by a smoothing spline with the smoothing parameter $\hat{\lambda}_f$, then the coverage property is lost, proving superiority of adaptive empirical Bayesian smoothing splines.

This paper is structured as follows. In Section 2 we describe the empirical Bayes approach and define some estimators. The asymptotic behaviour of the estimators is described in Section 3. In Section 4 we establish frequentist properties of a specific type of Bayesian credible set. In Section 5 we compare our approach for the selection of the smoothing parameter with generalised cross validation. Some numerical experiments can be found in Section 6. Section 7 contains some conclusions. We refer the reader to Section 8 for technical details and proofs.

2 Empirical Bayesian smoothing splines

Let us denote the minimiser of (2) by $\hat{f}_{\lambda,q}$, which is a natural smoothing spline of degree $2q - 1$ with knots at the observation points. This smoothing spline estimator is a linear estimator that satisfies for each λ and q

$$\hat{f}_{\lambda,q} = \hat{f}_{\lambda,q}(\mathbf{x}) = \mathbf{S}_{\lambda,q} \mathbf{Y}. \quad (3)$$

The positive-definite smoother matrix $\mathbf{S}_{\lambda,q} \in \mathbb{R}^{n \times n}$ is known explicitly and the vector $\mathbf{Y} = (Y_1, \dots, Y_n)^T$ comprises of the observations from model (1).

The L_2 -risk of the smoothing spline estimator (3) of $f \in \mathcal{W}_\beta$ is a function of λ and q known asymptotically; see Wahba (1990). In particular, with a suitable λ that minimises that risk, it holds for each q

$$\mathbb{E} \|\hat{f}_{\lambda,q} - f\|^2 \asymp n^{-\frac{2 \min(\beta,q)}{2 \min(\beta,q)+1}}. \quad (4)$$

This is of the order of the minimax risk for estimating $f \in \mathcal{W}_{\min(\beta, q)}$ in model (1). Apparently, to minimise this risk q should be larger than the maximum smoothness of f . This fact is known in a wider context of Tikhonov regularisation, see Lukas (1998) and references therein. However, we are not aware of any practical methods to select the optimal penalisation order q . In this work, estimates for both λ and q are obtained via the empirical Bayes method. We start by specifying a prior on the regression function f and on σ^2 .

Given (\mathbf{x}, λ, q) we endow σ^2 with an inverse-gamma prior with shape parameter a , and scale parameter b (both left unspecified for now), and given $(\sigma^2, \mathbf{x}, \lambda, q)$ we endow \mathbf{f} with a so called partially informative Gaussian prior with mean vector $\mathbf{0}$ and precision matrix $(\mathbf{S}_{\lambda, q}^{-1} - \mathbf{I}_n)/\sigma^2$, which we denote $PN\{\mathbf{0}, (\mathbf{S}_{\lambda, q}^{-1} - \mathbf{I}_n)/\sigma^2\}$ and admits a density which is proportional to

$$\left| \frac{\mathbf{S}_{\lambda, q}^{-1} - \mathbf{I}_n}{\sigma^2} \right|_+^{1/2} \exp \left\{ -\frac{1}{2\sigma^2} \mathbf{f}^T (\mathbf{S}_{\lambda, q}^{-1} - \mathbf{I}_n) \mathbf{f} \right\}, \quad (5)$$

where $|\cdot|_+$ represents the product of the non-zero eigenvalues (the smoother matrix $\mathbf{S}_{\lambda, q}$ has exactly q eigenvalues equal to 1; cf. Speckman 1985 and (31) in Section 8). This prior has two parts, a constant, non-informative prior on the null space of $\mathbf{S}_{\lambda, q}^{-1} - \mathbf{I}_n$, and a proper, degenerate Gaussian prior on the range of $\mathbf{S}_{\lambda, q}^{-1} - \mathbf{I}_n$; cf. Speckman and Sun, 2003.

We say that the prior on (\mathbf{f}, σ^2) is a partially-informative-Gaussian-inverse-gamma distribution, and denote it by

$$\Pi_{\lambda, q}(\cdot | \mathbf{x}) = PNIG\left\{\mathbf{0}, \mathbf{S}_{\lambda, q}^{-1} - \mathbf{I}_n, a, b\right\}. \quad (6)$$

This prior is conjugate for model (1) and the corresponding posterior distribution is

$$\Pi_{\lambda, q}(\cdot | \mathbf{Y}, \mathbf{x}) = PNIG\left\{\mathbf{S}_{\lambda, q} \mathbf{Y}, \mathbf{S}_{\lambda, q}^{-1}, a + \frac{n}{2}, b + \frac{1}{2} \mathbf{Y}^T (\mathbf{I}_n - \mathbf{S}_{\lambda, q}) \mathbf{Y}\right\}. \quad (7)$$

The posterior mean of (7) is $\hat{\mathbf{f}}_{\lambda, q}$, and the mean of the predictive posterior distribution can be shown to be the smoothing spline $\hat{f}_{\lambda, q}$. The prior (6) is improper but the corresponding posterior (7) is proper. By definition, the marginal posterior for σ^2 is an inverse-gamma distribution

$$\Pi_{\lambda, q}^{\sigma^2}(\cdot | \mathbf{Y}, \mathbf{x}) = \int \Pi_{\lambda, q}(d\mathbf{f}, \cdot | \mathbf{Y}, \mathbf{x}) = IG\left\{a + \frac{n}{2}, b + \frac{1}{2} \mathbf{Y}^T (\mathbf{I}_n - \mathbf{S}_{\lambda, q}) \mathbf{Y}\right\}, \quad (8)$$

and the marginal posterior for \mathbf{f} is a non-central, n -variate t-distribution (cf. Kotz and Nadarajah, 2004)

$$\Pi_{\lambda, q}^{\mathbf{f}}(\cdot | \mathbf{Y}, \mathbf{x}) = \int \Pi_{\lambda, q}(\cdot, d\sigma^2 | \mathbf{Y}, \mathbf{x}) = t_{2a+n-q}\left\{\mathbf{S}_{\lambda, q} \mathbf{Y}, \frac{2b + \mathbf{Y}^T (\mathbf{I}_n - \mathbf{S}_{\lambda, q}) \mathbf{Y}}{2a + n - q}, \mathbf{S}_{\lambda, q}\right\}. \quad (9)$$

The posterior distribution $\Pi_{\lambda, q}(\cdot | \mathbf{Y}, \mathbf{x})$ depends on λ and q , and on the (hyper-) parameters a and b . We select the unknown parameters λ and q in a data-driven way

via estimators $\hat{\lambda}$ and \hat{q} and plug them into the posterior (7) resulting in a new random measure, the empirical Bayes posterior, which is defined as

$$\Pi_{\hat{\lambda}, \hat{q}}(\cdot | \mathbf{Y}, \mathbf{x}) = \Pi_{\lambda, q}(\cdot | \mathbf{Y}, \mathbf{x})|_{(\lambda, q) = (\hat{\lambda}, \hat{q})}. \quad (10)$$

Following Robbins (1955), the empirical Bayes method consists of setting λ and q to maximisers of the marginal likelihood of model (1) under the Bayesian paradigm. We designate the mean of the empirical Bayes marginal posterior (10) as the *adaptive empirical Bayesian smoothing spline*.

Since the data $\mathbf{Y} | (f, \sigma^2, \mathbf{x})$ are distributed like a $N(\mathbf{f}, \sigma^2 \mathbf{I}_n)$ random vector, and we endow $\mathbf{f} | (\sigma^2, \mathbf{x}, \lambda, q)$ with a $PN\{\mathbf{0}, (\mathbf{S}_{\lambda, q}^{-1} - \mathbf{I}_n)/\sigma^2\}$ prior, then $\mathbf{Y} | (\sigma^2, \mathbf{x}, \lambda, q)$ is distributed like a $PN\{\mathbf{0}, (\mathbf{I}_n - \mathbf{S}_{\lambda, q})/\sigma^2\}$ random vector. The variance $\sigma^2 | (\mathbf{x}, \lambda, q)$ is endowed with an $IG(a, b)$ prior, such that $(\mathbf{Y}, \sigma^2) | (\mathbf{x}, \lambda, q)$ is by definition jointly distributed like a $PNIG\{\mathbf{0}, \mathbf{I}_n - \mathbf{S}_{\lambda, q}, a, b\}$ random vector. By integrating out σ^2 , $\mathbf{Y} | (\mathbf{x}, \lambda, q)$ is distributed like a $t_{2a-q}\{\mathbf{0}, 2b(\mathbf{I}_n - \mathbf{S}_{\lambda, q})^-/(2a - q)\}$ random vector, where the superscript “ $-$ ” indicates the pseudo-inverse. It can be shown that this distribution admits a density with respect to an appropriate dominating measure, resulting, up to some constant $h_n(a, b)$, in the following marginal log-likelihood for (λ, q) ,

$$\ell_n(\lambda, q | a, b) = -\left(a + \frac{n-q}{2}\right) \log \left\{ \mathbf{Y}^T (\mathbf{I}_n - \mathbf{S}_{\lambda, q}) \mathbf{Y} + 2b \right\} + \frac{1}{2} \log |\mathbf{I}_n - \mathbf{S}_{\lambda, q}|_+.$$

The hyper-parameters a and b do not play an important role in our approach so we set $a = q/2$, $b = 0$ with the convention that $0^0 = 1$ (this corresponds to placing an improper prior on σ^2). This does simplify the expressions that follow, but a and b can be set to any non-negative value that is $o(n)$ and does not depend on λ or q , without affecting our results. We obtain, up to a constant, the marginal log-likelihood

$$\ell_n(\lambda, q) = \ell_n(\lambda, q | q/2, 0) = -\frac{n}{2} \log \left\{ \mathbf{Y}^T (\mathbf{I}_n - \mathbf{S}_{\lambda, q}) \mathbf{Y} \right\} + \frac{1}{2} \log |\mathbf{I}_n - \mathbf{S}_{\lambda, q}|_+, \quad (11)$$

where (λ, q) lives on $(0, \infty) \times \mathbb{N}$; note that $h_n(q/2, 0)$ does not depend on λ or q .

The dependence of (11) on q is rather implicit, so that it is convenient to represent the smoother matrix as $\mathbf{S}_{\lambda, q} = \mathbf{\Phi} \{ \mathbf{I}_n + \lambda n \text{diag}(\boldsymbol{\eta}_q) \}^{-1} \mathbf{\Phi}^T$. Here $\mathbf{\Phi}$ is the Demmler-Reinsch basis matrix, such that $\mathbf{\Phi}^T \mathbf{\Phi} = \mathbf{\Phi} \mathbf{\Phi}^T = \mathbf{I}_n$, and $\boldsymbol{\eta}_q = (\eta_{q,1}, \dots, \eta_{q,n})^T$, see Section 8.1 for details. In particular, (31) in Section 8 gives an approximation of $\eta_{q,i}$ as a function of q . With this, we can re-express (11) as

$$\ell_n(\lambda, q) = -\frac{n}{2} \log \left(\sum_{i=q+1}^n \frac{X_i^2 \lambda n \eta_{q,i}}{1 + \lambda n \eta_{q,i}} \right) + \frac{1}{2} \sum_{i=q+1}^n \log \frac{\lambda n \eta_{q,i}}{1 + \lambda n \eta_{q,i}}, \quad (12)$$

where $\mathbf{X} = (X_1, \dots, X_n) = \mathbf{\Phi}^T \mathbf{Y}$. Further, based on the approximations from (31), $\ell_n(\lambda, q)$ is continuously differentiable.

A maximiser $(\hat{\lambda}, \hat{q})$ of $\ell_n(\lambda, q)$ is found by means of estimating equations, as zeroes of appropriately rescaled partial derivatives of $\ell_n(\lambda, q)$. Our estimating equations for λ and

q are respectively

$$\begin{aligned}
T_\lambda(\lambda, q) &= -\frac{2\lambda}{n^2} \left\{ \mathbf{Y}^T (\mathbf{I}_n - \mathbf{S}_{\lambda, q}) \mathbf{Y} \right\} \frac{\partial \ell_n(q, \lambda)}{\partial \lambda} \\
&= \frac{1}{n} \sum_{i=q+1}^n \frac{X_i^2 \lambda n \eta_{q,i}}{(1 + \lambda n \eta_{q,i})^2} - \frac{1}{n^2} \sum_{i=q+1}^n \frac{X_i^2 \lambda n \eta_{q,i}}{1 + \lambda n \eta_{q,i}} \sum_{i=q+1}^n \frac{1}{1 + \lambda n \eta_{q,i}}, \quad \text{and} \\
T_q(\lambda, q) &= -\frac{2q}{n^2} \left\{ \mathbf{Y}^T (\mathbf{I}_n - \mathbf{S}_{\lambda, q}) \mathbf{Y} \right\} \frac{\partial \ell_n(q, \lambda)}{\partial q} \\
&= \frac{1}{n} \sum_{i=q+1}^n \frac{X_i^2 \lambda n \eta_{q,i} \log(n \eta_{q,i})}{(1 + \lambda n \eta_{q,i})^2} - \frac{1}{n^2} \sum_{i=q+1}^n \frac{X_i^2 \lambda n \eta_{q,i}}{1 + \lambda n \eta_{q,i}} \sum_{i=q+1}^n \frac{\log(n \eta_{q,i})}{(1 + \lambda n \eta_{q,i})} \\
&\quad + \frac{1}{n} \sum_{i=q+1}^n \frac{\partial X_i^2}{\partial q} \frac{\lambda n \eta_i}{1 + \lambda n \eta_i},
\end{aligned} \tag{13}$$

See Section 8.1 for details on the derivation of these expressions. In particular, it is shown that the contribution of the last term in $T_q(\lambda, q)$ with $\partial X_i^2 / \partial q$ is negligible. Note that since $\ell_n(\lambda, q)$ is continuously differentiable, if $\hat{\lambda}_q$ solves $T_\lambda(\lambda, q) = 0$ for each q and \hat{q} solves $T_q(\hat{\lambda}_q, q) = 0$, then $(\hat{\lambda}, \hat{q}) = (\hat{\lambda}_{\hat{q}}, \hat{q})$. Note that for each q , $\hat{\lambda}_q$ is essentially the generalised maximum likelihood estimator from Wahba (1985); cf. (1.5) in Wahba (1985) and the criterion $T_\lambda(\lambda, q)$ above.

We briefly address some practical issues involving the optimisation of the criteria in (13). The estimate $\hat{\lambda}$ is taken on $[1/n, 1]$. Parameter q enters the eigenvalues $\eta_{q,i}$ according to (31) and X_i as the degree $2q - 1$ of basis Φ . While values $\eta_{q,i}$ are defined for each q , the degree of a spline is, in practice, typically an integer. One practical way to proceed in minimisation of (13) would be to restrict $q \in \mathbb{N}$. Alternatively, one could generalise splines to a fractional order (cf. Unser and Blu, 2000, for a representation of fractional splines in terms of fractional B-splines), which we do not pursue. Instead, we can use the fact that the contribution of the term with $\partial X_i^2 / \partial q$ is negligible (see Section 8.1 for details). Therefore, we suggest to relax $q \in (1/2, \log(n)]$ to be real-valued and in practice set $X_{q,i} = X_{[q],i}$, which allows to estimate non-integer q 's. Additionally, we also have to define $\mathcal{W}_\beta(M)$ for real-valued $\beta > 1/2$, which we do in Section 8, equation (29). Hence, throughout the paper both q and β are understood as real-valued numbers. In particular, all results and proofs hold also for $q, \beta \in \mathbb{N}$.

In practice, finding $(\hat{\lambda}, \hat{q})$ that optimises the criteria in (13) consists of finding $\hat{\lambda}_q$ that solves, for each q in some fine grid \mathbb{Q}_n , the criterion $T_\lambda(\lambda, q)$ up to an $o(1/n)$ factor, then finding $\hat{q} \in \mathbb{Q}_n$ that solves $T_q(\hat{\lambda}_q, q)$ and setting $(\hat{\lambda}, \hat{q}) = (\hat{\lambda}_{\hat{q}}, \hat{q})$. The grid $\mathbb{Q}_n = \{q_1, \dots, q_{N_n}\} \in (1/2, \log(n)]^{N_n}$ must be such that $|q_{i-1} - q_i| = o\{1/\log(n)\}$, $i = 1, \dots, N_n$, with $q_0 = 1/2$. This ensures that

$$n^{-\frac{2q_{i-1}}{2q_{i-1}+1}} = n^{-\frac{2q_i}{2q_i+1}} \{1 + o(1)\}, \quad i = 1, \dots, N_n,$$

which means that the discretisation is sufficiently fine.

3 Asymptotics of the solutions of the estimating equations

Fix some continuous regression function $f \in L_2$, and denote $\mathbf{B} = (B_1, \dots, B_n)^T = \Phi^T \mathbf{f}$ such that $\mathbb{E}\mathbf{X} = \mathbb{E}\Phi^T \mathbf{Y} = \mathbf{B}$. The oracle smoothing parameters will be defined as a solution to the system of equations

$$0 = \mathbb{E}T_\lambda(\lambda, q) = \frac{1}{n} \left\{ \sum_{i=q+1}^n \frac{B_i^2 \lambda n \eta_{q,i}}{(1 + \lambda n \eta_{q,i})^2} - \sum_{i=q+1}^n \frac{\sigma^2}{(1 + \lambda n \eta_{q,i})^2} + o(1) \right\}, \quad (14)$$

$$0 = \mathbb{E}T_q(\lambda, q) = \frac{1}{n} \sum_{i=q+1}^n \frac{B_i^2 \lambda n \eta_{q,i} \log(\lambda n \eta_{q,i})}{(1 + \lambda n \eta_{q,i})^2} + \log(1/\lambda) \mathbb{E}T_\lambda(\lambda, q), \quad (15)$$

where the expectation is taken under model (1). These expressions follow by several applications of Lemmas 1 and 2, similar to derivations in Krivobokova (2013). To solve this system assume that for each $q > 1/2$, equation (14) has a unique solution λ_q . Then equation (15) at $\lambda = \lambda_q$ becomes

$$0 = \mathbb{E}T_q(\lambda_q, q) = \frac{1}{n} \sum_{i=q+1}^n \frac{B_i^2 \lambda_q n \eta_{q,i} \log(\lambda_q n \eta_{q,i})}{(1 + \lambda_q n \eta_{q,i})^2}. \quad (16)$$

If this equation has a unique solution $\bar{\beta}$, then the solution to the system (14), (15) on $[1/n, 1] \times (1/2, \log(n)]$ will be called the oracle parameter $(\lambda_{\bar{\beta}}, \bar{\beta})$.

Apparently, the risk (4) depends on the relationship between q and β , whereby q should be chosen, while β is unknown. Therefore, we analyse both oracle parameters under two scenarios: a *low order penalty* scenario where $q \leq \max\{\beta > 1/2 : f \in \mathcal{W}_\beta(M)\}$, and a *high order penalty* scenario where $q > \max\{\beta > 1/2 : f \in \mathcal{W}_\beta(M)\}$. Here $\max\{\beta > 1/2 : f \in \mathcal{W}_\beta(M)\}$ is considered, since $f \in \mathcal{W}_\beta(M)$ does not preclude $f \in \mathcal{W}_{\beta'}(M)$ for some $\beta' > \beta$. Additionally to $f \in \mathcal{W}_\beta(M)$, we also discuss the case when f is an analytic signal. \mathcal{P}_∞ will denote the space of all analytic functions on $[0, 1]$ such that $\mathcal{P}_\infty \subset \mathcal{W}_\infty(M)$, while the space of all polynomials of degree $d - 1$ is denoted by \mathcal{P}_d , $d \in \mathbb{N}$.

3.1 Empirical Bayes estimate for λ

First we consider the solution to (14) for each $q > 1/2$.

Theorem 1. *Let $f \in \mathcal{W}_\beta(M)$, and assume that $\|f^{(\beta)}\|^2 > 0$. If $1/2 < q \leq \max\{\beta > 1/2 : f \in \mathcal{W}_\beta(M)\}$, then*

$$\lambda_q = \left[n \frac{\|f^{(q)}\|^2}{\sigma^2 \kappa_q(0, 2)} \{1 + o(1)\} \right]^{-2q/(2q+1)}, \quad (17)$$

where the constants $\kappa_q(m, l)$ are defined in Section 8.

If $q > \max\{\beta > 1/2 : f \in \mathcal{W}_\beta(M)\}$, then

$$\lambda_q \geq \left[n \frac{\|f^{(\beta)}\|^2}{\sigma^2 \kappa_q(0, 2)} \{1 + o(1)\} \right]^{-2q/(2\beta+1)}. \quad (18)$$

Moreover, for any $q > 1/2$, $\hat{\lambda}_q$ is consistent for λ_q and

$$\lambda_q^{-1/(4q)}(\hat{\lambda}_q/\lambda_q - 1) \xrightarrow{d} \mathcal{N}\left[0, \frac{2\kappa_q(2, 2)}{\{3\kappa_q(0, 2) - 2\kappa_q(0, 3)\}^2}\right], \quad \text{as } n \rightarrow \infty.$$

Proof of equations (17) and (18) follows from Lemmas 1 and 3 in Section 8. The consistency of $\hat{\lambda}_q$ and its asymptotic distribution in the case of $f \in \mathcal{W}_q(M)$ has been studied in Krivobokova (2013). Inspection of the proofs in Krivobokova (2013) shows that they hold with no changes for the case $q > \max\{\beta > 1/2 : f \in \mathcal{W}_\beta(M)\}$. Note that if $f \in \mathcal{P}_q$ such that $\|f^{(q)}\| = 0$, then $\lambda_q = \infty$.

3.2 Empirical Bayes estimate for q

First, consider the low penalty scenario where $1/2 < q \leq \max\{\beta > 1/2 : f \in \mathcal{W}_\beta(M)\}$ holds, so that in particular $f \in \mathcal{W}_q(M)$. By Lemma 3 (cf. Section 8),

$$\mathbb{E}T_q(\lambda_q, q) = -\lambda_q \log(1/\lambda_q) \|f^{(q)}\|^2 \{1 + o(1)\}. \quad (19)$$

Hence, for all $1/2 < q \leq \max\{\beta > 1/2 : f \in \mathcal{W}_\beta(M)\}$ the estimating equation $\mathbb{E}T_q(\lambda_q, q)$ remains strictly negative for $f \in \mathcal{W}_q(M) \setminus \mathcal{P}_q$ (that is as long as $\|f^{(q)}\| \neq 0$). If $f \in \mathcal{P}_q$, then $\mathbb{E}T_q(\lambda_q, q) = 0$.

Consider now the high order penalty scenario where $q > \max\{\beta > 1/2 : f \in \mathcal{W}_\beta(M)\}$. Contrary to the low penalty scenario, the sign of (16) is not characterised by just the assumption $f \in \mathcal{W}_\beta(M)$, which implies $f \notin \mathcal{W}_q(M)$. It turns out, that not every signal f that belongs to $\mathcal{W}_\beta(M)$ but not to $\mathcal{W}_{\beta+\delta}(M)$ for any $\delta > 0$ will be such that (16) is positive for $q > \max\{\beta > 1/2 : f \in \mathcal{W}_\beta(M)\}$. Such a mismatch between smoothness as “measured” by $\max\{\beta > 1/2 : f \in \mathcal{W}_\beta(M)\}$, and smoothness as “measured” by a change in the sign of the sum in (16) (which we can estimate), seems unavoidable (cf. Belitser and Enikeeva, 2008 for a similar issue in the context of hypothesis testing for smoothness in the Gaussian white noise model, and Giné and Nickl, 2010 in the context of Hölder smoothness in the construction of adaptive L_∞ credible bands in density estimation). From such issues stems, for example, the inability to construct adaptive credible sets in certain models with good coverage probability for $f \in \bigcup_{\beta \in B} \mathcal{W}_\beta$ if the range of smoothnesses B is large (cf. Low, 1997 and Section 4).

To “estimate” $\max\{\beta > 1/2 : f \in \mathcal{W}_\beta(M)\}$ for as large a family of models as possible, it is customary to remove from each model the functions for which the mismatch occurs, and thus consider the estimation problem over a smaller family. Possible functions to remove are those which are not *self-similar* (cf. Picard and Tribouley, 2000; Giné and Nickl, 2010), or that do not satisfy a *polished-tail* condition (cf. Szabó et al., 2014).

In our context, the set of functions with polished tails, call it $\mathcal{M} = \mathcal{M}(L, N, \rho)$, $L > 0$, $N \in \mathbb{N}$, $\rho \geq 2$, corresponds the set of all square integrable sequences such that

$$\frac{1}{n} \sum_{i=j}^n B_i^2 \leq \frac{L}{n} \sum_{i=j}^{\rho j} B_i^2, \quad N \leq j \leq n/\rho. \quad (20)$$

Such a condition has the role of excluding “irregular” signals, and unlike self-similarity conditions, it is not associated with any specific smoothness class. The condition ensures that the energy contained in the blocks $(B_j, \dots, B_{\rho j})$ does not surge over contributions of earlier blocks – the signal can only get more “polished” as one runs along the sequence B_i . This effectively excludes irregular signals that contain artefacts in their tails that influence the smoothness of the signal but are not detectable due to the noise.

For signals with polished tails, the criterium T_q can actually pick up on the smoothness of the signal. For fixed parameters L, N, ρ , there is a well defined notion of smoothness which we denote

$$\bar{\beta} = \max \{ \beta > 1/2 : f \in \mathcal{W}_\beta(M) \cap \mathcal{M} \}, \quad (21)$$

If f satisfies the polished tail conditions and is not in $\mathcal{W}_\beta(M)$, then by Lemma 3, for some $c > 0$,

$$\mathbb{E}T_q(\lambda_q, q) \geq c\lambda_q^{\beta/q} > 0, \quad (22)$$

for all large enough n . The conclusion is that for all sufficiently large n , $\mathbb{E}T_q(\lambda_q, q)$ changes signs at $\bar{\beta}$.

We conclude that the behaviour of $\mathbb{E}T_q(\lambda_q, q)$ can be described as follows. If for some $\beta > 1/2$, $f \in \mathcal{W}_\beta(M) \cap \mathcal{M}$ then the (continuous) criterion $\mathbb{E}T_q(\lambda_q, q)$ has a zero at $\min\{q \in \mathbb{Q}_n : q > \bar{\beta}\}$ since it is negative for $q \leq \bar{\beta}$, and positive for $q > \bar{\beta}$. If $f \in \mathcal{P}_\infty$ then $\mathbb{E}T_q(\lambda_q, q) \leq 0$, $q \in \mathbb{Q}_n$; if for some $d \in \mathbb{N}$, $f \in \mathcal{P}_d \setminus \mathcal{P}_{d-1}$ (such that f is a polynomial of degree exactly $d-1$), then $\mathbb{E}T_q(\lambda_q, q)$ is negative for $1/2 < q \leq d-1$ and is zero for $q \geq d$.

The proof of the following theorem is in Section 8.3.

Theorem 2. *Assume that for some $\beta > 1/2$, $f \in \mathcal{W}_\beta(M)$, then*

$$\mathbb{P}\{\beta < \hat{q} \leq \log(n)\} \rightarrow 1, \quad n \rightarrow \infty.$$

If furthermore $\beta = \bar{\beta}$ as defined in (21), and $f \in \mathcal{M}$, then

$$\hat{q} \xrightarrow{P} \bar{\beta}, \quad n \rightarrow \infty,$$

If for some $d \in \mathbb{N}$, $f \in \mathcal{P}_d$ then

$$\mathbb{P}\{d \leq \hat{q} \leq \log(n)\} \rightarrow 1, \quad n \rightarrow \infty.$$

As a side-note, the oracle λ_q can also be made more explicit for $f \in \mathcal{W}_\beta(M) \cap \mathcal{M}$. By (18) and (20), one finds that for each $f \in \mathcal{M}_{\bar{\beta}}(M) \cap \mathcal{M}$,

$$\lambda_q(f) = \left[n \frac{c_f M^2}{\sigma^2 \kappa_q(0, 2)} \{1 + o(1)\} \right]^{-2q/(2\bar{\beta}+1)}, \quad (23)$$

for some constant $0 < c_f \leq 1$, which depends on f and is bounded away from zero uniformly over $f \in \mathcal{W}_\beta(M) \cap \mathcal{M}$.

The exclusion of certain signals (so that the behaviour above holds) can be argued to be innocuous since the set of all signals that do not satisfy the polished-tail condition (or

that are not self-similar) is “small”. This can be justified following several arguments: removing such signals leaves the minimax rate (almost) unchanged so that the statistical problem does not become simpler; the probability that a function sampled from the prior does not satisfy such conditions is zero; there are also topological arguments for this. For a more extensive discussion cf. Giné and Nickl, 2010; Szabó et al., 2014 and the references therein. However, from the practical perspective, exactly which signals are removed is not relevant since one cannot check if the data come from a regression function satisfying a polished tail condition or not. In this sense one might as well implicitly exclude all signals for which (22) or (23) do not hold.

4 Bayesian credible sets as adaptive confidence sets

In Section 2 we propose a method for selecting the penalty order q of smoothing splines and the corresponding smoothing parameter λ_q . An immediate application of the consistency results in Section 3 is that \hat{q} and $\hat{\lambda}_{\hat{q}}$ can be directly plugged into the smoothing spline $\hat{f}_{\lambda,q}$ to obtain adaptive estimates for any continuous regression function $f \in L_2$. (This follows immediately from the consistency of the parameters, and standard arguments for smoothing spline estimators; cf. Wahba, 1990.) In this section we present another application: the construction of rate adaptive confidence sets based on the empirical Bayes posterior (10).

One of the often mentioned advantages of the Bayesian approach is that a posterior distribution provides statisticians with more than just point estimates. For appropriate priors, if the data are distributed according to a fixed distribution in the model, then with probability going to one, the posterior concentrates around this distribution. If this is the case, then for appropriate q and λ a small ℓ^2 -ball centred at the posterior mean $\hat{\mathbf{f}}_{\lambda,q}$, can capture most of the mass of the marginal posterior for \mathbf{f} . Since for each λ and q the posterior is known explicitly, simulating a high-probability region of the posterior (a credible set) and using it as a *frequentist* confidence set is of great appeal. However, it is known that such credible sets do not always have good frequentist coverage properties. This is more so the case when dealing with posteriors that adapt to the smoothness of the underlying signal to be estimated. In this section we adapt a technique developed by Szabó et al. (2014) for the Gaussian white noise model, to study the behaviour of a specific Bayesian credible set for our regression model (1). Complicating factors in our setup are that the variance of the noise is not assumed to be known, and that we work with two empirically chosen parameters (λ and q) simultaneously. We outline the technique in some detail since it is of independent interest.

We remind that the marginal posterior $\Pi_{\lambda,q}^{\mathbf{f}}(\cdot | \mathbf{Y}, \mathbf{x})$ equals

$$t_n\left(\hat{\mathbf{f}}_{\lambda,q}, \hat{\sigma}_{\lambda,q}^2 \mathbf{S}_{\lambda,q}\right), \quad \text{where} \quad \hat{\sigma}_{\lambda,q}^2 = \frac{1}{n} \mathbf{Y}^T (\mathbf{I} - \mathbf{S}_{\lambda,q}) \mathbf{Y}.$$

Representation properties of the multivariate t -distribution state that if \mathbf{f} is distributed like the marginal posterior above, then $\|\mathbf{f} - \hat{\mathbf{f}}_{\lambda,q}\|^2 / \hat{\sigma}_{\lambda,q}^2$ is distributed like $\mathbf{Z}^T \mathbf{S}_{\lambda,q} \mathbf{Z} / N$, where $\mathbf{Z} \sim N(\mathbf{0}, \mathbf{I}_n)$, $N \sim \chi_n^2$, and N is independent of \mathbf{Z} . Conclude that for any $\alpha \in (0, 1)$

there exists a (known, deterministic) sequence $r_n(\lambda, q)$ such that for every n, λ, q ,

$$\Pi_{\lambda, q}^{\mathbf{f}} \left(\|\mathbf{f} - \hat{\mathbf{f}}_{\lambda, q}\| \leq \hat{\sigma}_{\lambda, q} r_n(\lambda, q) \mid \mathbf{Y}, \mathbf{x} \right) = 1 - \alpha.$$

The level α is fixed for the remainder of this section. It is therefore natural to consider, for any $L \geq 1$, the empirical credible ball

$$\hat{\mathcal{C}}_n(L) = \left\{ \mathbf{f} \in \ell^2 : \|\mathbf{f} - \hat{\mathbf{f}}_{\hat{\lambda}_{\hat{q}}, \hat{q}}\| \leq \hat{\sigma}_{\hat{\lambda}_{\hat{q}}, \hat{q}} L r_n(\hat{\lambda}_{\hat{q}}, \hat{q}) \right\}. \quad (24)$$

By definition of the sequence $r_n(\lambda, q)$, for any $L \geq 1$,

$$\Pi_{\hat{\lambda}_{\hat{q}}, \hat{q}}^{\mathbf{f}} (\hat{\mathcal{C}}_n(L) \mid \mathbf{Y}, \mathbf{x}) \geq 1 - \alpha,$$

such that we can sample functions in $\hat{\mathcal{C}}_n(L)$ by sampling functions from the posterior and then keeping those that satisfy the inequality in (24). Such functions give a visual impression of the uncertainty in the point estimate $\hat{\mathbf{f}}_{\hat{\lambda}_{\hat{q}}, \hat{q}}$ – the adaptive empirical Bayesian smoothing spline. Note that since for each λ and q the posterior is known explicitly, simulating $\hat{\mathcal{C}}_n(L)$ is straightforward.

The following theorem is proved in Section 8.4 of Section 8.

Theorem 3. *Consider an interval $B = [\underline{b}, \bar{b}]$, where $1/2 < \underline{b} \leq \bar{b} < \infty$, and define $\mathcal{W} = \bigcup_{\beta \in B} \mathcal{W}_{\beta}(M) \cap \mathcal{M}(L, N, \rho)$. Then, for all large enough L ,*

$$\inf_{\mathbf{f} \in \mathcal{W}} \mathbb{P}_{\mathbf{f}} \{ \mathbf{f} \in \hat{\mathcal{C}}_n(L) \} = 1 + o(1), \quad \text{and} \quad (25)$$

$$\inf_{\mathbf{f} \in \mathcal{W}_{\beta}(M)} \mathbb{P}_{\mathbf{f}} \{ r_n(\hat{\lambda}_{\hat{q}}, \hat{q}) \leq K n^{-\beta/(2\beta+1)} \} = 1 + o(1), \quad \beta \in B, \quad (26)$$

for some large enough constant $K > 0$, depending on $L, M, \rho, \sigma^2, \underline{b}$, and \bar{b} .

Statement (25) is usually referred to as *honest coverage*, while (26) means that the credible ball $\hat{\mathcal{C}}_n(L)$ has a radius of the optimal order.

Ideally one would like to take $\mathcal{W} = \ell^2$. However, as mentioned in Section 3.2, it is known (cf. Low, 1997) that it is in general not possible to fulfil conditions (25) and (26) simultaneously if $\bar{b}/\underline{b} > 2$ and $\mathcal{W} = \ell^2$. To allow (25) and (26) to hold simultaneously for a wide (but bounded) range of smoothnesses one usually identifies “problematic” functions which are either removed from the model (cf. Giné and Nickl, 2010; Szabó et al., 2014) or replaced with a collection of so called surrogates, “non-problematic” replacement functions that retain the main features of interest of the functions that were removed from the model (cf. Genovese and Wasserman, 2008). Imposing an upper bound \bar{b} on the smoothness is also necessary if we are to establish (25) and (26). Such a bound can also be justified from a computational standpoint.

The constant L , which is the multiplicative factor for the radius of the credible set $\hat{\mathcal{C}}_n$, must be taken appropriately large for (25) and (26) to hold. It is possible to provide

an explicit lower-bound L by inspecting the constants in the proof of Theorem 3: for all sufficiently large n we may take

$$L \geq 1 + \{\kappa_q(0, 2)/\kappa_q(0, 1)\}^{1/2} = 1 + \{(2q - 1)/(2q)\}^{1/2},$$

so that uniformly over q , $L \geq 2$. Inspection of the computations in Section 8 shows that the level α appears associated with lower order terms so that L does not depend on α , even if we were to allow L to depend on q . Because of this one does not get exactly coverage $1 - \alpha$ and the credible sets $\hat{\mathcal{C}}_n$ are always conservative (in that the asymptotic probability of coverage is $1 > 1 - \alpha$).

It follows, in fact, from the inequalities established in the proof of Theorem 3 that the posterior distribution contracts (with respect to $\|\cdot\|$) around the true regression function at the optimal rate for $f \in \mathcal{W}_\beta$, when f is indeed in this space. To establish this, condition (20) is not needed.

5 Comparison with frequentist smoothing splines

In the frequentist framework there are several competing ways of selecting the smoothing parameter λ (for a fixed q). Typically, λ is selected as a minimiser of some asymptotically unbiased estimator of the risk $\mathbb{E}\|\hat{\mathbf{f}}_{\lambda,q} - \mathbf{f}\|^2$. Generalised cross validation, Mallow's C_p and Akaike's information criterion are particular examples of such estimates; let $\hat{\lambda}_f$ denote a minimiser of one of such criteria. If the regression function f belongs to $\mathcal{W}_\beta(M)$, and we set $q \in [\beta/2, \beta]$ such that $\beta/q \in [1, 2]$, then $\hat{\lambda}_f$ adapts. This means that generalised cross-validated smoothing splines adapt to the unknown smoothness β in the sense that the estimator $\hat{\lambda}_f$ is consistent for the oracle

$$\lambda_f \geq \left[n \frac{\|f^{(\beta)}\|^2}{\sigma^2 \kappa_q(1, 2)} \{1 + o(1)\} \right]^{-2q/(2\beta+1)}. \quad (27)$$

Furthermore, Theorem 3 of Krivobokova (2013) states that

$$\lambda_f^{-1/(4q)} (\hat{\lambda}_f / \lambda_f - 1) \xrightarrow{d} \mathcal{N} \left[0, \frac{2\kappa_q(4, 2)}{\{4\kappa_q(1, 2) - 3\kappa_q(1, 3)\}^2} \right], \quad \text{as } n \rightarrow \infty.$$

so that the asymptotic variance above can be much larger than that associated with the empirical Bayes estimate $\hat{\lambda}_\beta$ for the range of values of q for which the GCV $\hat{\lambda}_f$ adapts, see Krivobokova (2013) for more discussion and simulations.

In this section we investigate how the asymptotic variance of $\hat{\lambda}_f$ compares to that of $\hat{\lambda}_{\hat{q}}$; or, more specifically, we compare the distances of $\hat{\mathbf{f}}_{\hat{\lambda}_f, q}$ and of $\hat{\mathbf{f}}_{\hat{\lambda}_{\hat{q}}, \hat{q}}$ to the true regression function. We use the credible sets from the previous section as a proxy for this comparison: we bound the probability that the regression function belongs to a ball centered at $\hat{\mathbf{f}}_{\hat{\lambda}_f, q}$ with a radius that assures coverage for the Bayesian credible set $\hat{\mathcal{C}}_n(L)$ – by construction,

that is the radius of $\hat{C}_n(1)$, since we are only interested here in coverage, and not honest coverage.

The proof of the following theorem can be found in Section 8.5.

Theorem 4. *Assume that the regression function f belongs to $\mathcal{W}_\beta(M)$, $\beta > 1/2$, such that with probability going to 1 the radius of the credible set $\hat{C}_n(1)$ is $\sigma r_n(\lambda_\beta, \beta)$. Define*

$$\hat{\mathcal{D}}_n = \hat{\mathcal{D}}_n(\hat{\mathbf{f}}_{\hat{\lambda}_f, q}) = \left\{ \mathbf{f} \in \ell^2 : \|\mathbf{f} - \hat{\mathbf{f}}_{\hat{\lambda}_f, q}\| \leq \sigma r_n(\lambda_\beta, \beta) \right\}.$$

Then, for any q such that $\hat{\lambda}_f$ adapts to the smoothness β , i.e., for any $q \in [\beta/2, \beta]$,

$$\mathbb{P}_f \left\{ \mathbf{f} \in \hat{\mathcal{D}}_n(\hat{\mathbf{f}}_{\hat{\lambda}_f, q}) \right\} = o(1), \quad n \rightarrow \infty.$$

The conclusion is that $\hat{\lambda}_f$ can somewhat adapt to the smoothness of the regression function, but at the cost of a high asymptotic variance. Using the fact that for each fixed q , the empirical Bayes selected $\hat{\lambda}_q$ has much lower asymptotic variance, we show that the smoothing parameter $\hat{\lambda}_{\hat{q}}$ outperforms $\hat{\lambda}_f$: if the centre of the empirical Bayes credible ball (the adaptive empirical Bayesian smoothing spline) is replaced by a frequentist smoothing spline with the smoothing parameter $\hat{\lambda}_f$, then the coverage property is lost. Note that even if this were not the case, the empirical Bayes smoothing spline would still adapt to a wider range of smoothnesses than the risk-based smoothing spline (which adapts only within a $[q, 2q]$ range).

6 Numerical simulations

The following simulation study aims to verify our theoretical findings in finite samples and compare frequentist and proposed adaptive empirical Bayesian smoothing splines in terms of the average mean squared error.

In all settings the Monte-Carlo sample $M = 1000$, the sample size is $n = 1000$, the design points are fixed and equidistant $x = i/n$, $i = 1, \dots, n$ and $\sigma = 0.1^2$. The results for other sample sizes and higher σ s were found conceptually similar and are not reported. We consider two mean functions f_j , $j = 1, 2$ that are scaled by its range. Function f_1 , shown on the top left plot of Figure 1, has a known decay of its Demmler-Reinsch coefficients, namely we set

$$f_1(x) = \sum_{i=q+1}^n \psi_{\beta, i}(x) (i+1)^{-\beta} \cos(2i), \quad \beta = 3,$$

where $\psi_{\beta, i}$ is the i th basis function of the Demmler-Reinsch basis of degree $\beta = 3$, defined in (28). The second function is analytical $f_2(x) = \cos(5\pi x)$ (Figure 2) with an exponential decay of its Demmler-Reinsch coefficients.

We estimated both functions by smoothing splines with the empirical Bayesian smoothing parameter $\hat{\lambda}_{\hat{q}}$ and with the smoothing parameter $\hat{\lambda}_f$. To get $\hat{\lambda}_{\hat{q}}$, first $\hat{\lambda}_q$ is obtained for $q = 1, 2, \dots, 6$ as a solution to $T_\lambda(\lambda, q) = 0$ and then \hat{q} is set to the nearest integer to q^* such that $T_q(\hat{\lambda}_q, q^*) = 0$. The smoothing parameter $\hat{\lambda}_f$ has been calculated by generalized

cross-validation for different values of $q = 2, \dots, 6$.

We compare the resulting smoothing parameter estimators $\hat{\lambda}_{\hat{q}}$ and $\hat{\lambda}_f$ for $q = 2, \dots, 6$ in terms of the sample mean and sample variance in the Monte Carlo sample (see $\text{mean}(\hat{\lambda})$ and $\text{var}(\hat{\lambda})$ in Table 1). Further, we compare the empirical average mean squared error (AMSE) of the resulting estimators, that is

$$A(\hat{\lambda}) = \frac{1}{Mn} \sum_{i=1}^M \sum_{j=1}^n \left\{ \hat{f}_i(x_j, \hat{\lambda}) - f(x_j) \right\}^2,$$

where \hat{f}_i denotes a smoothing spline estimator in the i th Monte Carlo simulation. The values of $A(\hat{\lambda})$ are given in Table 1 together with the AMSE ratio, which is defined as $R = A(\hat{\lambda}_f)/A(\hat{\lambda}_{\hat{q}})$, so that values of $R > 1$ imply superiority of the adaptive empirical Bayesian estimator.

Let us first consider the simulation results for f_1 . According to the theoretical results on $\hat{\lambda}_f$ discussed in Section 5, as long as $q \geq \beta/2$, its oracle $\lambda_f \geq \text{const } n^{-2q/(2\beta+1)}$, leading to $\mathbb{E}\|\hat{\mathbf{f}}(\hat{\lambda}_f) - \mathbf{f}\|^2 \asymp n^{-2\beta/(2\beta+1)}$. For $\beta = 3$ we expect to observe that λ_f for $q = 2$ is larger than for $q = 3$ and goes very fast to zero for $q > \beta = 3$. Moreover, $\hat{f}(\hat{\lambda})$ for $\hat{\lambda}_{\hat{q}}$ and $\hat{\lambda}_f$ for $q = 2, \dots, 6$ should all have the same convergence rates and differ only in constants. Results of the simulations given in Table 1 confirm these theoretical findings. We observe also that the means of $\hat{\lambda}_{\hat{q}}$ and $\hat{\lambda}_f$ for $q = 3$ are of the same rate, but the variance of $\hat{\lambda}_f$ is much higher, which is also visible in the boxplots, given on the bottom left plot of Figure 1. The differences between $A(\hat{\lambda})$, shown on the bottom right plot of Figure 1, are less pronounced, but from the AMSE ratio given in Table 1 we find that the empirical Bayesian smoothing spline estimator outperforms the frequentist smoothing spline uniformly in q . The smallest difference is observed between $A(\hat{\lambda}_{\hat{q}})$ and $A(\hat{\lambda}_f)$ for the true $q = \beta = 3$, which is, of course, unknown in practice and is not estimated in the frequentist framework. Finally, we remark that the empirical Bayesian estimator of q appeared to be very reliable for this value of β :

	EB	GCV				
		$q = 2$	$q = 3$	$q = 4$	$q = 5$	$q = 6$
f_1						
$\text{mean}(\hat{\lambda})$	$5.7 \cdot 10^{-12}$	$1.8 \cdot 10^{-08}$	$7.1 \cdot 10^{-12}$	$1.8 \cdot 10^{-15}$	$4.5 \cdot 10^{-19}$	$1.2 \cdot 10^{-22}$
$\text{var}(\hat{\lambda})$	$5.0 \cdot 10^{-25}$	$2.3 \cdot 10^{-17}$	$1.2 \cdot 10^{-23}$	$1.7 \cdot 10^{-30}$	$2.0 \cdot 10^{-37}$	$2.4 \cdot 10^{-44}$
$A(\hat{\lambda})$	$2.5 \cdot 10^{-06}$	$2.9 \cdot 10^{-06}$	$2.6 \cdot 10^{-06}$	$2.7 \cdot 10^{-06}$	$2.9 \cdot 10^{-06}$	$2.9 \cdot 10^{-06}$
R	—	1.147912	1.032582	1.075522	1.119951	1.139055
f_2						
$\text{mean}(\hat{\lambda})$	$1.7 \cdot 10^{-19}$	$5.9 \cdot 10^{-09}$	$1.6 \cdot 10^{-11}$	$3.4 \cdot 10^{-14}$	$3.6 \cdot 10^{-17}$	$2.8 \cdot 10^{-19}$
$\text{var}(\hat{\lambda})$	$1.1 \cdot 10^{-36}$	$2.0 \cdot 10^{-18}$	$2.4 \cdot 10^{-23}$	$4.8 \cdot 10^{-28}$	$4.0 \cdot 10^{-34}$	$2.4 \cdot 10^{-38}$
$A(\hat{\lambda})$	$1.5 \cdot 10^{-06}$	$3.8 \cdot 10^{-06}$	$2.1 \cdot 10^{-06}$	$1.9 \cdot 10^{-06}$	$1.8 \cdot 10^{-06}$	$1.6 \cdot 10^{-06}$
R	—	2.483676	1.406107	1.266803	1.164454	1.040263

Table 1: Simulation results for functions f_1 and f_2 .

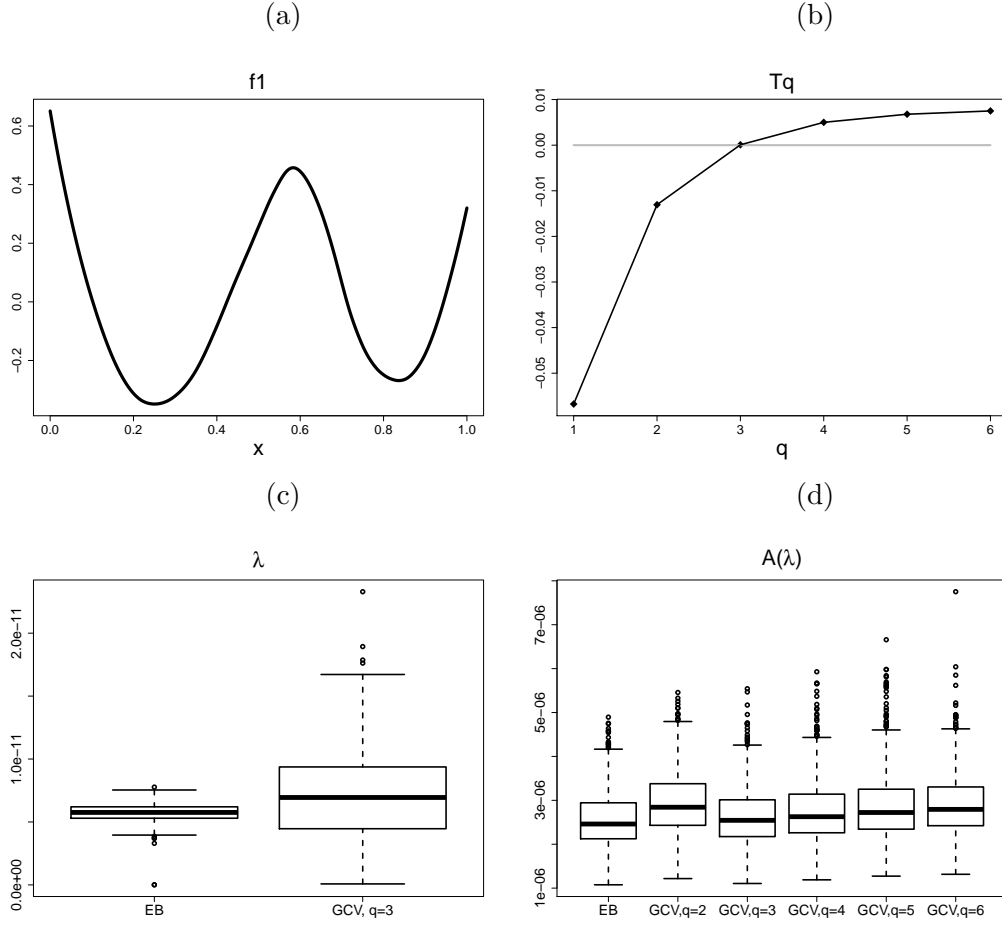


Figure 1: (a) Function f_1 , (b) Criterion $T_q(\hat{\lambda}_q, q)$ for f_1 , (c) Boxplots of $\hat{\lambda}_q$ and $\hat{\lambda}_f$ obtained with GCV and $q = \beta = 3$, (d) Boxplots of $A(\hat{\lambda}_q)$ (EB) and $A(\hat{\lambda}_f)$

out of $M = 1000$ samples in 998 cases $\hat{q} = 3$ has been obtained and in two cases $\hat{q} = 4$. The estimating equation $T_q(\hat{\lambda}_q, q)$ is shown on the top right plot of Figure 1.

Now we consider the simulation results for the analytical function $f_2(x) = \cos(5\pi x)$ which is known to have exponentially decaying Demmler-Reinsch coefficients. In this setting, criterion $T_q(\hat{\lambda}_q, q)$ should remain negative, asymptotically approaching zero from below, which is visible in the top right plot of Figure 2. We estimated $\hat{\lambda}_q$ setting $\hat{q} = 6$ if $T_q(\hat{\lambda}_q, q) < 0$ for $q = 6$. For function f_2 estimator \hat{q} is slightly more variable, resulting in $\hat{q} = 6$ in 986 cases and in $\hat{q} = 5$ in 14 cases. This can be due to the fact that in small samples it is difficult to distinguish between an exponential decay $\exp(-\pi i)$ and a decay with $(\pi i)^{-q}$, $q > 5$, $i = 1, \dots, n$. Also, $\hat{\lambda}_f$ for $q = 2, \dots, 6$ with generalized cross-validation have been calculated. It appears that $\hat{\lambda}_q$ and $\hat{\lambda}_f$ with $q = 6$ have the same rate, but again, $\hat{\lambda}_f$ is much more variable. For these smoothing parameters we also observe, that

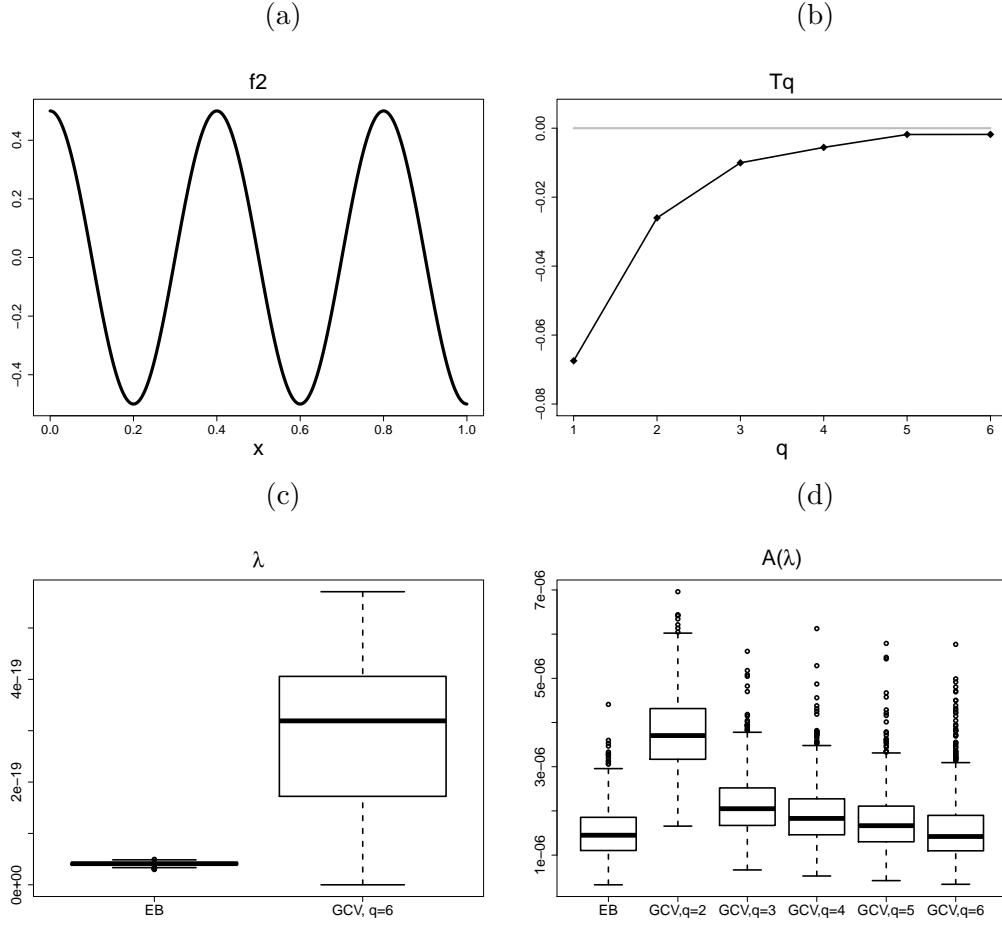


Figure 2: (a) Function f_2 , (b) Criterion $T_q(\hat{\lambda}_q, q)$ for f_2 , (c) Boxplots of $\hat{\lambda}_{\hat{q}}$ and $\hat{\lambda}_f$ obtained with GCV and $q = 6$, (d) Boxplots of $A(\hat{\lambda}_{\hat{q}})$ (EB) and $A(\hat{\lambda}_f)$

the corresponding AMSE ratio is closest to one. In general, for this function f_2 the adaptive empirical Bayesian smoothing spline estimator again outperforms frequentist splines uniformly in q with the largest AMSE ratio of about 2.5 for $q = 2$.

Finally, we remark on the implementation of the procedure. It is well-known that the spline based basis with knots at observations becomes numerically unstable for higher q s. In fact, it seems impossible to get numerically stable Demmler-Reinsch basis for the natural spline space for $q > 3$ with usual approaches. Instead, we relied on an approximation based on the Demmler-Reinsch basis (28) for \mathcal{W}_q . The details of this approach will be reported elsewhere, but the implementation in R is available from the authors on request.

7 Conclusions

The selection of the order of smoothing splines in non-parametric regression is a topic mostly absent from the literature. The empirical Bayes method is shown to provide an adequate framework to produce data driven choices for this parameter. Although the dependence of the prior on the parameter q – which controls the order of the smoothing spline – is rather implicit, if the regression function has a well defined smoothness (determined with the help of so the so called polished-tail condition), then \hat{q} is consistent and we identify the smallest Sobolev space containing the regression function. Hence, our adaptive empirical Bayesian smoothing spline estimator (which is the mean of the empirical Bayes posterior) adapts to the underlying smoothness of the signal.

High probability regions of the empirical posterior are shown to have good frequentist coverage properties. For a large class of functions the size of these regions is shown to adapt to the underlying smoothness of the signal, effectively quantifying the amount of uncertainty of the empirical Bayesian smoothing spline estimator. These results are used to show that frequentist smoothing splines are outperformed by empirical Bayesian smoothing splines.

Acknowledgements

We would like to thank the editor, the associate editor and both referees for the extremely useful remarks, that helped to improve the paper substantially. We express also our gratitude to Francisco Rosales for his implementation of the Demmler-Reinsch basis and helpful discussions.

8 Technical results and proofs

8.1 Demmler-Reinsch basis and estimating equations

Let $\{\psi_i\}_{i=1}^\infty$ denote the Demmler-Reinsch basis of $\mathcal{W}_\beta(M)$, such that it holds

$$\int_0^1 \psi_{\beta,i}(x) \psi_{\beta,j}(x) dx = \delta_{ij} = \nu_{\beta,i}^{-1} \int_0^1 \psi_{\beta,i}^{(\beta)}(x) \psi_{\beta,j}^{(\beta)}(x) dx. \quad (28)$$

Hence, any $f \in \mathcal{W}_\beta(M)$ can be represented as $f = \sum_{i=1}^\infty \theta_{\beta,i} \psi_{\beta,i}$ for $\theta_{\beta,i} = \int_0^1 f(x) \psi_{\beta,i}(x) dx$, and $\|f^{(\beta)}\|^2 = \sum_{i=1}^\infty \theta_{\beta,i}^2 \nu_{\beta,i} < M^2$. We also define Sobolev spaces for real-valued β s via

$$\mathcal{W}_\beta(M) = \{f : f \in \mathcal{C}^{[\beta]-1}[0, 1], \|f^{(\beta)}\|^2 = \sum_{i=1}^\infty \theta_{\beta,i}^2 \nu_{\beta,i}^{\beta/[\beta]} < M^2\}. \quad (29)$$

The Demmler-Reinsch basis of the natural spline space of degree $2q - 1$ with knots at observations $\mathcal{S}_{2q-1}(\mathbf{x})$ is defined via

$$\sum_{k=1}^n \phi_{q,i}(x_k) \phi_{q,j}(x_k) = \delta_{ij} = \eta_{q,i}^{-1} \int_0^1 \phi_{q,i}^{(q)}(x) \phi_{q,j}^{(q)}(x) dx. \quad (30)$$

With the Demmler-Reinsch basis the smoother matrix can be represented as $\mathbf{S}_{\lambda,q} = \mathbf{\Phi}_q \{ \mathbf{I}_n + \lambda n \text{diag}(\boldsymbol{\eta}_q) \}^{-1} \mathbf{\Phi}_q^T$, where $\mathbf{\Phi}_q = \mathbf{\Phi}_q(\mathbf{x}) = [\phi_{q,1}(\mathbf{x}), \dots, \phi_{q,n}(\mathbf{x})] = [\phi_{q,j}(x_i)]_{i,j=1}^n$. For the eigenvalues $\boldsymbol{\eta}_q$, Theorem 2.2. in Speckman (1985) establishes the following approximation

$$\eta_{q,1} = 0, \dots, \eta_{q,q} = 0, \quad n\eta_{q,i} = \pi^{2q}(i-q)^{2q}\{1+o(1)\}, \quad i = q+1, \dots, n, \quad (31)$$

where the $o(1)$ term is uniform over $i = o\{n^{2/(2q+1)}\}$, as $n \rightarrow \infty$.

For $q = \beta$ in Speckman (1985) is shown that $\sqrt{n}\phi_{\beta,i} = \psi_{\beta,i}\{1+o(1)\}$, $n\eta_{\beta,i} = \nu_{\beta,i}\{1+o(1)\}$, $n^{-1}B_{\beta,i}^2 = \theta_{\beta,i}^2\{1+o(1)\}$, as well as $n^{-1}\sum_{i=\beta+1}^n B_{\beta,i}^2 n\eta_{\beta,i} = \|f^{(\beta)}\|^2\{1+o(1)\}$, where $B_{\beta,i} = \sum_{j=1}^n \phi_{\beta,j}(x_i)f(x_j)$. In case $\beta \neq q$, that is $f \in \mathcal{W}_\beta(M)$, but the Demmler-Reinsch basis of degree q is used, we find in a similar fashion

$$\frac{1}{n} \sum_{i=q+1}^n B_{q,i}^2 (n\eta_{q,i})^{\min(\beta,q)/q} = \|f^{(\min(\beta,q))}\|^2\{1+o(1)\}.$$

The estimating equation $T_\lambda(\lambda, q)$ defined in Section 2 is derived in the same way as those in Krivobokova (2013) using the fact that $d\mathbf{S}_{\lambda,q}/d\lambda = -\lambda^{-1}(\mathbf{S}_{\lambda,q} - \mathbf{S}_{\lambda,q}^2)$; $T_q(\lambda, q)$ follows in the same way with a few extra remarks. Firstly, to determine the derivatives of eigenvalues $\eta_{q,i}$ with respect to q we employ representation (31). Secondly, we use that $\sum_{i=1}^n X_{q,i}^2 = \mathbf{Y}^t \mathbf{\Phi}_q \mathbf{\Phi}_q^T \mathbf{Y} = \sum_{i=1}^n Y_i^2$, which is independent of q and implies that the contribution of $\partial X_i^2/\partial q$ is negligible. This can be seen as follows. For $j = j(n) \propto \lambda^{-\alpha/(2q)}$ with any $\alpha \in (0, 1)$, so that $\lambda n\eta_j = o(1)$, similar to the proof of Lemma 2 in Krivobokova (2013) follows

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n \frac{\partial X_{q,i}^2}{\partial q} \frac{\lambda n\eta_{q,i}}{1 + \lambda n\eta_{q,i}} &= \frac{1}{n} \sum_{i=1}^n \frac{\partial X_{q,i}^2}{\partial q} - \frac{1}{n} \sum_{i=1}^n \frac{\partial X_{q,i}^2}{\partial q} \frac{1}{1 + \lambda n\eta_{q,i}} \\ &= \frac{1}{n} \sum_{i=j+1}^n \frac{\partial X_{q,i}^2}{\partial q} \frac{\lambda n\eta_{q,i}}{1 + \lambda n\eta_{q,i}} \{1 + o(1)\}, \end{aligned}$$

that is, $\partial X_{q,i}^2/\partial q$ contributes only to the tail part of the sum. Next, observe that $X_{q,i}/\sqrt{n} = B_{q,i}/\sqrt{n} + O_p(\sigma/\sqrt{n})$, so that

$$\frac{1}{n} \sum_{i=j+1}^n \frac{\partial X_{q,i}^2}{\partial q} \frac{\lambda n\eta_{q,i}}{1 + \lambda n\eta_{q,i}} = \frac{2\sigma}{\sqrt{n}} \sum_{i=j+1}^n \frac{\partial B_{q,i}/\sqrt{n}}{\partial q} \{O_p(1) + B_{q,i}/\sigma\} \frac{\lambda n\eta_{q,i}}{1 + \lambda n\eta_{q,i}}.$$

Since $n^{-1}\sum_{i=1}^n B_{q,i}^2 n\eta_{q,i} < \infty$, we can bound $|B_{q,i}|/\sqrt{n} \leq \text{const}(i-q)^{-(2q+1+\epsilon)/2}$ for some $\epsilon > 0$ and for $i > j$ it follows that $B_{q,i} = O(1)$ for $\lambda = O(n^{-2q/\{\alpha(2q+1+\epsilon)\}})$, which includes λ_q for a suitable choice of $\alpha \in (0, 1)$. Since $(\lambda n\eta_{q,i})^\delta/(1 + \lambda n\eta_{q,i}) \leq 1$, for some $\delta \in [0, 1]$, showing that $|\partial B_{q,i}/\partial q| = O\{\log(n)\}|B_{q,i}|$ will imply that

$$\begin{aligned} \left| \frac{1}{n} \sum_{i=j+1}^n \frac{\partial X_{q,i}^2}{\partial q} \frac{\lambda n\eta_{q,i}}{1 + \lambda n\eta_{q,i}} \right| &\leq \frac{2\sigma \log(n)}{\sqrt{n}} \sum_{i=j+1}^n \left| \frac{B_{q,i}}{\sqrt{n}} \right| (\lambda n\eta_{q,i})^{\min(\beta,q)/2q} O_p(1) \\ &= \frac{\lambda^{\min(\beta,q)/q} \log(n)}{\sqrt{\lambda^{\min(\beta,q)/q} n}} o_p(1). \end{aligned}$$

To see that $|\partial B_{q,i}/\partial q| = O\{\log(n)\}|B_{q,i}|$, note that for $i > q$

$$B_{q,i} = \sum_{j=1}^n f(x_j) \phi_{q,i}(x_j) = \sum_{j=1}^n f(x_j) \sum_{k < j} \alpha_{i,k} (x_j - x_k)^{2q-1},$$

for some suitable $\alpha_{i,k}$, $k < j$, since $\phi_i(x_j)$ is a spline function of degree $2q - 1$, which is orthogonal to a polynomial space. Taking the derivative with respect to q , gives the result. We conclude that the contribution of $\partial X_{q,i}^2/\partial q$ is negligible and therefore subindices q in $X_{q,i}$ and $B_{q,i}$ will be subsequently omitted.

8.2 Auxiliary lemmas

Lemma 1. *Let $\mathbf{S}_{\lambda,q}$ be the smoother matrix. Then, for all $q > 1/2$, $l \in \mathbb{N}$, $m \in \mathbb{N} \cup \{0\}$, and λ such that $\lambda \rightarrow 0$, $n\lambda \rightarrow \infty$, as $n \rightarrow \infty$,*

$$\text{tr} \{(\mathbf{I}_n - \mathbf{S}_{\lambda,q})^m \mathbf{S}_{\lambda,q}^l\} = \sum_{i=1}^n \frac{(\lambda n \eta_{q,i})^m}{(1 + \lambda n \eta_{q,i})^{m+l}} = \lambda^{-1/(2q)} \kappa_q(m, l) \{1 + o(1)\},$$

where the $o(1)$ term depends only on m and l , and the constants $\kappa_q(m, l)$ are

$$\kappa_q(m, l) = \frac{\Gamma\{m + 1/(2q)\} \Gamma\{l - 1/(2q)\}}{2\pi q \Gamma(l + m)}. \quad (32)$$

Proof. See Krivobokova (2013). □

Lemma 2. *Consider $\eta_{q,i}$ as in (31). Then, as $n \rightarrow \infty$ such that $\lambda \rightarrow 0$ for all $m \in \mathbb{N} \cup \{0\}$, $r, l \in \mathbb{N}$, $q > 1/2$*

$$\sum_{i=q+1}^n \frac{(\lambda n \eta_{q,i})^m \{\log(n \eta_{q,i})\}^r}{(1 + \lambda n \eta_{q,i})^{m+l}} = \lambda^{-1/(2q)} \{\log(1/\lambda)\}^r \kappa_q(m, l) \{1 + o(1)\}.$$

Proof. Follows the same lines as Lemma 1. □

Lemma 3. *Let $m \in \mathbb{N}$, $q > 1/2$, and $\lambda \rightarrow 0$ such that $\lambda n \rightarrow \infty$. If $f \in \mathcal{W}_q$,*

$$\frac{1}{n} \sum_{i=q+1}^n \frac{B_i^2 \lambda n \eta_{q,i} \log(\lambda n \eta_{q,i})}{(1 + \lambda n \eta_{q,i})^m} = -\lambda \log(1/\lambda) \|f^{(q)}\|^2 \{1 + o(1)\}. \quad (33)$$

If $f \in \mathcal{W}_\beta \cap \mathcal{M}$, but $f \notin \mathcal{W}_{\beta+\epsilon}$, $\epsilon > 0$, and $q > \beta > 1/2$, then

$$\frac{1}{n} \sum_{i=q+1}^n \frac{B_i^2 \lambda n \eta_{q,i} \log(\lambda n \eta_{q,i})}{(1 + \lambda n \eta_{q,i})^m} \geq c \lambda^{\beta/q}, \quad (34)$$

$$\frac{1}{n} \sum_{i=q+1}^n \frac{B_i^2 \lambda n \eta_{q,i}}{(1 + \lambda n \eta_{q,i})^m} \geq c \lambda^{\beta/q}, \quad (35)$$

for some sufficiently small constant $c > 0$. If f is a polynomial of degree $d - 1$ (such that $f \in \mathcal{P}_d$ and $f \notin \mathcal{P}_{d-1}$) and $q \geq d$ then

$$\frac{1}{n} \sum_{i=q+1}^n \frac{B_i^2 \lambda n \eta_{q,i} \log(\lambda n \eta_{q,i})}{(1 + \lambda n \eta_{q,i})^m} = 0. \quad (36)$$

Proof. We remind that if $f \in \mathcal{W}_q$ then

$$\frac{1}{n} \sum_{i=q+1}^n B_i^2 n \eta_{q,i} = \|f^{(q)}\|^2 \{1 + o(1)\} < \infty.$$

Obviously, if f is a polynomial of degree $d - 1$, $d \in \mathbb{N}$ ($f \in \mathcal{P}_d$, $f \notin \mathcal{P}_{d-1}$), then $\|f^{(q)}\| = 0$, $q \geq d$ and (36) follows.

To prove (33) let $j = j_n(\alpha)$ be like in the proof of Lemma 2 in Krivobokova (2013); more specifically, take $j = \pi^{-1} \lambda^{-\alpha/(2q)}$ for $\alpha \in (0, 1)$, so that $\lambda n \eta_{q,j} = o(1)$; following the proof of that lemma, for each $\alpha \in (0, 1)$,

$$\frac{1}{n} \sum_{i=q+1}^n \frac{B_i^2 \lambda n \eta_{q,i} \log(\lambda n \eta_{q,i})}{(1 + \lambda n \eta_{q,i})^m} = -\frac{\lambda}{n} \left[\sum_{i=q+1}^j B_i^2 n \eta_{q,i} \log \frac{1}{\lambda n \eta_{q,i}} \{1 + o(1)\} - O\left(\sum_{i=j+1}^n B_i^2 n \eta_{q,i} \right) \right].$$

If $f \in \mathcal{W}_q$ such that $n^{-1} \sum_{i=q+1}^n B_i^2 n \eta_{q,i} = \|f^{(q)}\|^2 \{1 + o(1)\} < \infty$, whence for any choice of $j \rightarrow \infty$ the second term on the right-hand-side is negligible compared to the first. Conclude that for each $\alpha \in (0, 1)$,

$$\lambda \log \lambda \|f^{(q)}\|^2 \{1 + o(1)\} \leq \frac{1}{n} \sum_{i=q+1}^n \frac{B_i^2 \lambda n \eta_{q,i} \log(\lambda n \eta_{q,i})}{(1 + \lambda n \eta_{q,i})^m} \leq (1 - \alpha) \lambda \log \lambda \|f^{(q)}\|^2 \{1 + o(1)\},$$

(for an $o(1)$ term which is uniform over α) such that (33) follows by taking $\alpha \rightarrow 0$.

Finally, we establish (34). Let $f \in \mathcal{W}_\beta(M)$, $f \notin \mathcal{W}_{\beta+\epsilon}$, $\epsilon > 0$, $q > \beta + \epsilon$. Consider the function

$$g_r(x) = \frac{x^r \log(x)}{(1+x)^m}, \quad r \in (0, 1), \quad x > 0.$$

This continuous function is negative on the interval $(0, 1)$ and positive on $(1, \infty)$. It converges to zero when $x \rightarrow 0$, has a global minimum at some point $m_r \in (0, 1)$, is zero at $x = 1$, has a global maximum at $M_r \in (1, \infty)$, and converges monotonically to zero as $x \rightarrow \infty$. Let $j_{-2} \leq j_{-1} \leq j_0 \leq j_1 \leq j_2$ be sequences of integers depending on n such that: $g_{1-\beta/q}(\lambda n \eta_{q,j_{-2}}) = o(1)$ and $g_{1-\beta/q}(\lambda n \eta_{q,j_{-1}}) = o(1)$ with $\lambda n \eta_{q,j_{-2}} < m_{1-\beta/q} < \lambda n \eta_{q,j_{-1}}$ and $j_{-2} \rightarrow \infty$; $\lambda n \eta_{q,j_0} < 1 \leq \lambda n \eta_{q,j_0+1}$; and for some (small) constant $0 < c < g_1(M_1)$, $g_1(\lambda n \eta_{q,i}) \geq c$, $i = j_1, \dots, j_2$, such that $g_1(\lambda n \eta_{q,j_1-1}) < c$, $g_1(\lambda n \eta_{q,j_2+1}) < c$. (Note that m_r , M_r , and the sequences j_k , $k = -2, \dots, 2$ depend on r and m .)

Splitting the summation along these sequences, and using the bounds above

$$\begin{aligned} \frac{1}{n} \sum_{i=q+1}^{j_0} \frac{B_i^2 \lambda n \eta_{q,i} \log(\lambda n \eta_{q,i})}{(1 + \lambda n \eta_{q,i})^m} &\geq \\ &\geq -\lambda^{\beta/q} \frac{1}{n} \left\{ o(1) \sum_{i \in I} B_i^2 (n \eta_{q,i})^{\beta/q} - g_{1-\beta/q}(m_{1-\beta/q}) \sum_{i=j_{-2}+1}^{j_{-1}} B_i^2 (\lambda n \eta_{q,i})^{\beta/q} \right\}, \end{aligned}$$

where $I = \{q+1, \dots, j_{-2}\} \cup \{j_{-1}, \dots, j_0\}$. Since by definition $f \in \mathcal{W}_\beta$, then we have $n^{-1} \sum_{i=q+1}^n B_i^2 (\lambda n \eta_{q,i})^{\beta/q} = \|f^{(\beta)}\|^2 \{1 + o(1)\} < \infty$, so, since $j_{-2} \rightarrow \infty$ we conclude that the lower bound in the previous display is $-o(\lambda^{\beta/q}) \|f^{(\beta)}\|^2$, which lower bounds the sum of the negative terms. As for the sum of the positive terms,

$$\frac{1}{n} \sum_{i=j_0+1}^n \frac{B_i^2 \lambda n \eta_{q,i} \log(\lambda n \eta_{q,i})}{(1 + \lambda n \eta_{q,i})^m} \geq \frac{1}{n} \sum_{i=j_1}^{j_2} \frac{B_i^2 \lambda n \eta_{q,i} \log(\lambda n \eta_{q,i})}{(1 + \lambda n \eta_{q,i})^m} \geq \frac{c}{n} \sum_{i=j_1}^{j_2} B_i^2.$$

Simple computations show we can take $O(\lambda^{-1/(2q)}) = j_1 \leq j_2/\rho$, for any $\rho > 2$. By the polished-tail condition the previous display can further be lower bounded by

$$\frac{c}{n} \sum_{i=j_1}^{\rho j_1} B_i^2 \geq \frac{c (n \eta_{q, \rho j_1})^{-\frac{\beta+\epsilon}{q}}}{n} \sum_{i=j_1}^{\rho j_1} B_i^2 (n \eta_{q,i})^{\frac{\beta+\epsilon}{q}} \geq \frac{O(j_1^{-2(\beta+\epsilon)})}{nL} \sum_{i=j_1}^n B_i^2 (n \eta_{q,i})^{\frac{\beta+\epsilon}{q}}.$$

The sum $\sum_{i=j_1}^n B_i^2 (n \eta_{q,i})^{\frac{\beta+\epsilon}{q}}/n$ converges to infinity as long as $j_1 = o(n)$, $\epsilon > 0$. The power $j_1^{-2(\beta+\epsilon)}$ is $o(\lambda^{\beta/q})$ for every $\epsilon > 0$ but it is $O(\lambda^{\beta/q})$ if $\epsilon = 0$. Conclude that for small enough ϵ , namely such that $j_1^{-2(\beta+\epsilon)}$ is of a larger order than the $o(\lambda^{\beta/q})$ term from the bound on the negative terms, the sum of the positive terms dominates. In fact, picking an appropriate sequence $\epsilon \rightarrow 0$ one can show that the sum of the positive terms is lower bounded by a small enough multiple of $\lambda^{\beta/q}$. The bound in (34) follows. The bound in (35) follows in the same way by adjusting the definition of the sequences j_1 and j_2 . \square

Lemma 4 (Lemma 3 from Krivobokova, 2013). *Let $\lambda \rightarrow 0$ be such that $n\lambda \rightarrow \infty$. Then, for any $l \in \mathbb{N}$,*

$$\begin{aligned} \mathbf{f}^T (\mathbf{I}_n - \mathbf{S}_{\lambda,q}) \mathbf{S}_{\lambda,q}^{1+l} \mathbf{f} |_{\lambda=\lambda_q} &= \sigma^2 \text{tr}(\mathbf{S}_{\lambda,q}^2) |_{\lambda=\lambda_q} \{1 + o(1)\}, \\ \mathbf{f}^T (\mathbf{I}_n - \mathbf{S}_{\lambda,q})^{1+l} \mathbf{S}_{\lambda,q} \mathbf{f} |_{\lambda=\lambda_q} &= o(\lambda_q^{-1/(2q)}). \end{aligned}$$

Lemma 5. *Let $T_{C_p}(\lambda)$ be the estimating equation for λ_f as defined in Krivobokova (2013), λ_f the solution to $\mathbb{E}T_{C_p}(\lambda) = 0$, and $\hat{\lambda}_f$ the solution to $T_{C_p}(\lambda) = 0$. Then*

$$\mathbb{E} \left\{ \lambda_f^{-1/(4q)} (\hat{\lambda}_f / \lambda_f - 1) \right\}^2 \leq \frac{4\kappa_q(4, 2)}{\{4\kappa_q(1, 2) - 3\kappa_q(1, 3)\}^2} \{1 + o(1)\},$$

such that in particular $\mathbb{V}(\hat{\lambda}_f / \lambda_f - 1) = o(1)$.

Proof. By definition of $\hat{\lambda}_f$, for some $\tilde{\lambda}$ between $\hat{\lambda}_f$ and λ_f ,

$$0 = T_{C_p}(\hat{\lambda}_f) = T_{C_p}(\lambda_f) + T'_{C_p}(\tilde{\lambda})(\hat{\lambda}_f - \lambda_f).$$

It is known (cf. the supplementary materials to Krivobokova, 2013) that,

$$\begin{aligned} \mathbb{V}\{T_{C_p}(\lambda_f)\} &= 2\frac{\sigma^4}{n^2}\kappa_q(4, 2)\lambda_f^{-1/(2q)}\{1 + o(1)\}, \\ \mathbb{E}\{T'_{C_p}(\tau\lambda_f)\} &= \mathbb{E}\{T'_{C_p}(\lambda_f)\}\left\{1 + \frac{2q-1}{2q}O(\tau-1)\right\} \\ &= \frac{\sigma^2}{n}\lambda_f^{-1-1/(2q)}(4\kappa_q(1, 2) - 3\kappa_q(1, 3))\left\{1 + \frac{2q-1}{2q}O(\tau-1)\right\}, \end{aligned}$$

for $\tau \in [1 - \epsilon, 1 + \epsilon]$ and small enough $\epsilon > 0$. Let $A(\epsilon)$ be the event $\{|\hat{\lambda}_f/\lambda_f - 1| < \epsilon\}$, and write $\tilde{\lambda} = \tau\lambda_f$, so that

$$\begin{aligned} 2\frac{\sigma^4}{n^2}\kappa_q(4, 2)\lambda_f^{-1/(2q)}\{1 + o(1)\} &= \\ &= \mathbb{E}\left\{(\hat{\lambda}_f - \lambda_f)^2 \mathbb{E}[T'_{C_p}(\tilde{\lambda})^2 | \hat{\lambda}_f]\right\} \geq \mathbb{E}\left\{(\hat{\lambda}_f - \lambda_f)^2 \mathbb{E}[T'_{C_p}(\tau\lambda_f)^2 | \hat{\lambda}_f, A(\epsilon)] \mathbb{P}(A(\epsilon))\right\} \\ &= \mathbb{E}\left\{(\hat{\lambda}_f - \lambda_f)^2 \mathbb{E}[T'_{C_p}(\tau\lambda_f) | A(\epsilon)]^2 \{1 + o(1)\}\right\} \\ &= \mathbb{E}\left\{(\hat{\lambda}_f - \lambda_f)^2 \mathbb{E}[T'_{C_p}(\lambda_f)]^2 \mathbb{E}\left[1 + \frac{2q-1}{2q}O(\tau-1) | A(\epsilon)\right]^2 \{1 + o(1)\}\right\}, \end{aligned}$$

where we use the fact that by definition $\mathbb{E}\{T_{C_p}(\lambda_f)\} = 0$, and Jensen's inequality. Since the previous display holds for all small enough ϵ , set ϵ so that $1 + (2q-1)O(\tau-1)/(2q) \geq 1/\sqrt{2}$, say, whence

$$\mathbb{E}\{\hat{\lambda}_f/\lambda_f - 1\}^2 \leq \lambda_f^{1/(2q)} \frac{4\kappa_q(4, 2)}{\{4\kappa_q(1, 2) - 3\kappa_q(1, 3)\}^2} \{1 + o(1)\} = o(1).$$

□

8.3 Proof of Theorem 2

Assume $f \in L_2$. First, note that for some $\tilde{\lambda}$ between $\hat{\lambda}_q$ and λ_q , a.s.,

$$T_q(\hat{\lambda}_q, q) = T_q(\lambda_q, q) + (\hat{\lambda}_q - \lambda_q) \frac{\partial T_q(\lambda, q)}{\partial \lambda} \Big|_{\lambda=\tilde{\lambda}}. \quad (37)$$

We have

$$\begin{aligned} \mathbb{V}T_q(\lambda_q, q) &= \frac{2\sigma^2}{q^2 n^2} \left[\frac{\{\log(1/\lambda_q)\}^2}{\lambda_q^{1/(2q)}/\kappa_q(2, 2)} \sigma^2 \{1 + o(1)\} + 2 \sum_{i=q+1}^n B_i^2 \left\{ \frac{\lambda_q n \eta_{q,i} \log(n \eta_{q,i})}{(1 + \lambda_q n \eta_{q,i})^2} \right\}^2 \right] \\ &= O\left\{ \lambda_q^{-1/(2q)} \log(1/\lambda_q)^2 n^{-2} \right\} \rightarrow 0, \quad n \rightarrow \infty. \end{aligned}$$

Hence, $T_q(\lambda_q, q) - \mathbb{E}T_q(\lambda_q, q) \xrightarrow{P} 0$ for any $q > 1/2$. In fact, for any $\epsilon > 0$, and \mathbb{Q}_n as defined at the end of Section 2,

$$\begin{aligned} \mathbb{P}\left(\max_{q \in \mathbb{Q}_n} |T_q(\lambda_q, q) - \mathbb{E}T_q(\lambda_q, q)| > \epsilon\right) &\leq \mathbb{P}\left(\sum_{q \in \mathbb{Q}_n} |T_q(\lambda_q, q) - \mathbb{E}T_q(\lambda_q, q)| > \epsilon\right) \leq \\ &\sum_{q \in \mathbb{Q}_n} \mathbb{P}\left(|T_q(\lambda_q, q) - \mathbb{E}T_q(\lambda_q, q)| > \epsilon/|\mathbb{Q}_n|\right) \leq \frac{|\mathbb{Q}_n|^3 w}{\epsilon^2} \max_{q \in \mathbb{Q}_n} \mathbb{V}T_q(\lambda_q, q), \end{aligned}$$

by Chebyshev's inequality. We assume without loss of generality that $|\mathbb{Q}_n| = O\{1/\log(n)^\alpha\}$, for some $\alpha > 2$ such that the upper-bound in the previous display converges to zero, implying that $T_q(\lambda_q, q) - \mathbb{E}T_q(\lambda_q, q) \xrightarrow{P} 0$ takes place uniformly over $q \in \mathbb{Q}_n$. Note also that $\mathbb{V}T_q(\lambda_q, q) = o[\{\mathbb{E}T_q(\lambda_q, q)\}^2]$. The remaining term in (37) is controlled using the same arguments as Lemma 5 so we omit the details.

Consider a polished tail $f \in \mathcal{L}_2$. Without loss of generality say that T_q , $q \notin \mathbb{Q}_n$, is obtained by linear interpolation of the values of T_q at \mathbb{Q}_n . From the discussion in Sections 3.1 and 3.2, the criterion $\mathbb{E}T_q(\lambda_q, q)$ is strictly negative for $q \leq \bar{\beta}$ and is positive for $q > \bar{\beta}$, whence

$$\begin{aligned} \mathbb{P}\left\{\max_{q \in \{r \in \mathbb{Q}_n: r < \bar{\beta}\}} T_q(\hat{\lambda}_q, q) < 0, \min_{q \in \{r \in \mathbb{Q}_n: r > \bar{\beta}\}} T_q(\hat{\lambda}_q, q) > 0\right\} &\leq \\ &\leq \mathbb{P}\left\{\sup_{q \in (1/2, \bar{\beta})} T_q(\hat{\lambda}_q, q) < 0, \inf_{q \in (\bar{\beta}, \log(n)]} T_q(\hat{\lambda}_q, q) > 0\right\} \leq \mathbb{P}\{\hat{q} = \bar{\beta}\}, \end{aligned}$$

The first inequality follows from the fact that we linearly interpolate the criterion, and the second inequality follows from the definition of \hat{q} . From the computations above, the left-hand-side of the previous display converges to 1 as $n \rightarrow \infty$, so that the second statement of the theorem follows. For $f \in \mathcal{W}_\beta(M)$, $\beta > 1/2$, by the same arguments as above,

$$\begin{aligned} \mathbb{P}\left[\max_{q \in \{r \in \mathbb{Q}_n: r < \beta\}} T_q(\hat{\lambda}_q, q) < 0, \min_{q \in \{r \in \mathbb{Q}_n: r > \beta\}} T_q(\hat{\lambda}_q, q) = o_p\{T_q(\hat{\lambda}_\beta, \beta)\}\right] &\leq \\ &\leq \mathbb{P}\left[\sup_{q \in (1/2, \beta)} T_q(\hat{\lambda}_q, q) < 0, \inf_{q \in (\beta, \log(n)]} T_q(\hat{\lambda}_q, q) = o_p\{T_q(\hat{\lambda}_\beta, \beta)\}\right] \leq \mathbb{P}[\beta \leq \hat{q} \leq \log(n)], \end{aligned}$$

and the first statement of the theorem follows since the left-hand-side of the previous display converges to 1, as $n \rightarrow \infty$. The last statement of the theorem follows by using the exact same argument.

8.4 Proof of Theorem 3

If we assume that for some countable set $Q_n \in \mathbb{Q}_n$,

$$\inf_{\mathbf{f} \in \mathcal{W}} \mathbb{P}_f(\hat{q} \in Q_n) = 1 + o(1), \quad (38)$$

then (25) follows if we show that uniformly over $q \in Q_n$,

$$\inf_{\mathbf{f} \in \mathcal{W}} \mathbb{P}_f\{\|\mathbf{f} - \hat{\mathbf{f}}_{\hat{\lambda}_q, q}\| \leq \hat{\sigma}_{\hat{\lambda}_q, q} L r_n(\hat{\lambda}_q, q)\} = 1 + o(1/|Q_n|). \quad (39)$$

For each $q \in Q_n$, $f \in \mathcal{W}$, consider intervals $L_q = [\lambda_q, \bar{\lambda}_q]$ – in fact $L_q(f) = [\lambda_q(f), \bar{\lambda}_q(f)]$ – with $\lambda_q \in L_q$ and $\bar{\lambda}_q/\lambda_q \rightarrow 1$, such that

$$\inf_{\mathbf{f} \in \mathcal{W}} \mathbb{P}_f\{\hat{\lambda}_q \in L_q(f)\} = 1 + o(1). \quad (40)$$

If there exist constants C_1, C_2, C_3, C_4 , such that uniformly over $f \in \mathcal{W}$ and $q \in Q_n$,

$$\inf_{\lambda \in L_q} r_n^2(\lambda, q) \geq C_1^2 \frac{\bar{\lambda}_q^{-1/(2q)}}{n}, \quad (41)$$

$$\inf_{f \in \mathcal{W}} \mathbb{P}_f\left(\inf_{\lambda \in L_q} \hat{\sigma}_{\lambda, q}^2 \geq C_2^2\right) = 1 + o(1), \quad (42)$$

$$\sup_{\lambda \in L_q} \|\mathbb{E}_f \hat{\mathbf{f}}_{\lambda, q} - \mathbf{f}\|^2 \leq C_3^2 \frac{\bar{\lambda}_q^{-1/(2q)}}{n}, \quad (43)$$

$$\inf_{f \in \mathcal{W}} \mathbb{P}_f\left(\sup_{\lambda \in L_q} \|\hat{\mathbf{f}}_{\lambda, q} - \mathbb{E}_f \hat{\mathbf{f}}_{\lambda, q}\|^2 \leq C_4^2 \bar{\lambda}_q^{-1/(2q)}/n\right) = o(1/|Q_n|), \quad (44)$$

then, if L is large enough such that $C_4 + C_3 \leq L C_2 C_1$, conclude

$$\inf_{f \in \mathcal{W}} \mathbb{P}_f\left\{\sup_{\lambda \in L_q} \|\hat{\mathbf{f}}_{\lambda, q} - \mathbb{E}_f \hat{\mathbf{f}}_{\lambda, q}\| + \sup_{\lambda \in L_q} \|\mathbb{E}_f \hat{\mathbf{f}}_{\lambda, q} - \mathbf{f}\| \leq \inf_{\lambda \in L_q} \hat{\sigma}_{\lambda, q} L r_n(\lambda, q)\right\} = 1 + o(1/|Q_n|),$$

such that by the triangle inequality, (39) holds.

If (38) holds, then (26) follows if we show that uniformly over $q \in Q_n$,

$$\inf_{f \in \mathcal{W}_{\beta}(M)} \mathbb{P}_f\{r_n(\hat{\lambda}_q, q) \leq K n^{-\beta/(2\beta+1)}\} = 1 + o(1/|Q_n|). \quad (45)$$

Since the risk of the smoothing spline satisfies (4), it is only possible to establish (45) if $q \geq \beta$. Consider $s_n(\beta) = \max\{q \in Q_n : q \leq \beta\}$. By the first and second statements of Theorem 2 we conclude that if we take $Q_n = Q_n \cap (s_n, \bar{\beta} + 1)$, then $\mathbb{P}(\hat{q} \in Q_n) = 1 + o(1)$, such that every $q \in Q_n$ is such that $q \geq \beta$. Note also that Q_n can be taken such that $|Q_n| = o\{\log(n)^2\}$, say. These facts are used when checking (42), (44), and (45) below.

Condition (40) follows from the consistency results in Section 3 so that it suffices to check conditions (41)–(45), which we do in the remainder of this section.

8.4.1 Checking condition (41)

The sequence $r_n(\lambda, q)$ is implicitly defined by

$$\mathbb{P}\{R_n(\lambda, q) \leq r_n(\lambda, q)^2\} = 1 - \alpha, \quad \text{with} \quad R_n(\lambda, q) \sim \frac{1}{N} \mathbf{Z}^T \mathbf{S}_{\lambda, q} \mathbf{Z} \sim \frac{1}{N} \sum_{i=q+1}^n \frac{\epsilon_i^2}{1 + \lambda n \eta_{q, i}},$$

where $\mathbf{Z} \sim N(\mathbf{0}, \mathbf{I}_n)$ and $N \sim \mathcal{X}_n^2$, with \mathbf{Z} and N independent, and $\boldsymbol{\epsilon} = \boldsymbol{\Phi}^T \mathbf{Z} \sim N(\mathbf{0}, \mathbf{I}_n)$. Since $R_n(\lambda, q)$ decreases with λ , the infimum in (41) is attained at $\lambda = \bar{\lambda}_q$. Set $R =$

$R_n(\bar{\lambda}_q, q)$, $r = r_n(\bar{\lambda}_q, q)$ and note that by the one-sided Chebychev's inequality, $r^2 \geq \mathbb{E}R - \{(1 - \alpha)/\alpha\}^{1/2}(\mathbb{V}R)^{1/2}$. Further, since N^{-1} is independent of ϵ , and has an inverse-chi-squared distribution with n degrees of freedom such that $\mathbb{E}N^{-1} = (n - 2)^{-1}$,

$$\mathbb{E}R = \mathbb{E} \frac{1}{N} \sum_{i=q+1}^n \frac{\mathbb{E}\epsilon_i^2}{1 + \bar{\lambda}_q n \eta_{q,i}} = \frac{1}{n-2} \text{tr}(\mathbf{S}_{\bar{\lambda}_q, q}) = \frac{\bar{\lambda}_q^{-1/(2q)}}{n-2} \kappa_q(0, 1) \{1 + o(1)\}.$$

By similar computations, since $\mathbb{V}N^{-1} = 2(n - 2)^{-2}(n - 4)^{-1}$,

$$\mathbb{V}R = 3 \left(\mathbb{E} \frac{1}{N} \right)^2 \text{tr}(\mathbf{S}_{\bar{\lambda}_q, q}^2) + \mathbb{V} \frac{1}{N} \text{tr}(\mathbf{S}_{\bar{\lambda}_q, q})^2 + 3 \mathbb{V} \frac{1}{N} \text{tr}(\mathbf{S}_{\bar{\lambda}_q, q}^2) = o\left\{(\mathbb{E}R)^2\right\}.$$

Conclude that (41) holds for $C_1^2 < \kappa_q(0, 1)\{1 + o(1)\}$.

8.4.2 Checking condition (42)

We may take $L_q = [c\bar{\lambda}_q, \bar{\lambda}_q/c]$, for some $c = 1 + o(1) < 1$. For $\lambda \in L_q$, we bound, a.s.,

$$\hat{\sigma}_{\lambda, q}^2 = \frac{1}{n} \mathbf{Y}^T (\mathbf{I}_n - \mathbf{S}_{\lambda, q}) \mathbf{Y} \geq \frac{c^2}{n} \sum_{i=q+1}^n \frac{X_i^2 \bar{\lambda}_q n \eta_{q,i}}{1 + \bar{\lambda}_q n \eta_{q,i}} \geq \frac{c^2}{n} \mathbf{Y}^T (\mathbf{I}_n - \mathbf{S}_{\bar{\lambda}_q, q}) \mathbf{Y}.$$

Denote this lower-bound by $S = S_n(\bar{\lambda}_q, q, c)$. By Chebyshev's inequality,

$$\mathbb{P}\left(\inf_{\lambda \in L_q} \hat{\sigma}_{\lambda, q}^2 \geq C_2^2\right) \geq 1 - \mathbb{P}\left\{S \leq \mathbb{E}S - (\mathbb{E}S - C_2^2)\right\} \geq 1 - \frac{\mathbb{V}S}{\mathbb{V}S + (\mathbb{E}S - C_2^2)^2}.$$

Using Lemma 1 we have

$$\mathbb{E}S = \frac{c^2}{n} \left\{ \sigma^2 \text{tr}(\mathbf{I}_n - \mathbf{S}_{\bar{\lambda}_q, q}) + \mathbf{f}^T (\mathbf{I}_n - \mathbf{S}_{\bar{\lambda}_q, q}) \mathbf{f} \right\} \geq c^2 \sigma^2 - c^2 \sigma^2 \frac{\text{tr}(\mathbf{S}_{\bar{\lambda}_q, q})}{n} = \sigma^2 \{1 + o(1)\}.$$

By Lemma 1, and the bound on $\mathbf{f}^T (\mathbf{I}_n - \mathbf{S}_{\bar{\lambda}_q, q})^2 \mathbf{f}$ from the next section,

$$\mathbb{V}S = \frac{1}{n^2} \left[2\sigma^4 \text{tr}\{(\mathbf{I}_n - \mathbf{S}_{\bar{\lambda}_q, q})^2\} + 4\sigma^2 \mathbf{f}^T (\mathbf{I}_n - \mathbf{S}_{\bar{\lambda}_q, q})^2 \mathbf{f} \right] = O\left(\frac{\mathbb{E}S}{n}\right).$$

Condition (42) holds by taking $C_2^2 < \sigma^2 \{1 + o(1)\}$. Since the bound in previous display is $o\{1/\log(n)^2\}$ we further assume that the bound holds uniformly over $q \in Q_n$.

8.4.3 Checking condition (43)

Let L_q be as in the previous section. For $\lambda \in L_q$ the squared bias term satisfies

$$\|\mathbb{E}\hat{\mathbf{f}}_{\lambda, q} - \mathbf{f}\|^2 = \frac{1}{n} \sum_{i=q+1}^n \frac{B_i^2(\lambda n \eta_{q,i})^2}{(1 + \lambda n \eta_{q,i})^2} \leq \frac{1}{n} \sum_{i=q+1}^n \frac{B_i^2(\bar{\lambda}_q n \eta_{q,i})^2}{(1 + \bar{\lambda}_q n \eta_{q,i})^2} \leq \frac{1}{c^2} \|(\mathbf{S}_{\lambda_q, q} - \mathbf{I}_n) \mathbf{f}\|^2.$$

By Lemma 4 and simple manipulations,

$$\|(\mathbf{S}_{\lambda_q, q} - \mathbf{I}_n) \mathbf{f}\|^2 = \frac{1}{n} \mathbf{f}^T (\mathbf{S}_{\lambda_q, q} - \mathbf{I}_n) \mathbf{S}_{\lambda_q, q} \mathbf{f} - \frac{1}{n} \mathbf{f}^T (\mathbf{S}_{\lambda_q, q} - \mathbf{I}_n) \mathbf{f} \leq \sigma^2 \frac{\text{tr}(\mathbf{S}_{\lambda_q, q})}{n} \{1 + o(1)\}.$$

The statement follows by Lemma 1, for $C_3^2 > \sigma^2 \kappa_q(0, 1)\{1 + o(1)\}$.

8.4.4 Checking condition (44)

For L_q as in the previous sections, uniformly over $\lambda \in L_q$ the variance term satisfies, a.s.,

$$\|\hat{\mathbf{f}}_{\lambda,q} - \mathbb{E}\hat{\mathbf{f}}_{\lambda,q}\|^2 = \|\mathbf{S}_{\lambda,q}(\mathbf{Y} - \mathbf{f})\|^2 = \frac{1}{n} \sum_{i=q+1}^n \frac{(X_i - B_i)^2}{(1 + \lambda n \eta_{q,i})^2} \leq \frac{1}{n} \sum_{i=q+1}^n \frac{(X_i - B_i)^2}{(1 + \underline{\lambda}_q n \eta_{q,i})^2},$$

where $\mathbf{X} - \mathbf{B} = \mathbf{\Phi}^T(\mathbf{Y} - \mathbf{f}) \sim N(\mathbf{0}, \sigma^2 \mathbf{I}_n)$. Denote the upper-bound in the previous display by $V = V_n(\underline{\lambda}_q, q)$. By Lemma 1, the expectation of V is

$$\mathbb{E}V = \frac{\sigma^2}{n} \text{tr}(\mathbf{S}_{\underline{\lambda}_q,q}^2) = \frac{\sigma^2}{n} \underline{\lambda}_q^{-1/(2q)} \kappa_q(0, 2) \{1 + o(1)\} = \sigma^2 c^{-1/q} \kappa_q(0, 2) \frac{\bar{\lambda}_q^{-1/(2q)}}{n} \{1 + o(1)\};$$

the respective variance is $\mathbb{V}V = 3\sigma^2 n^{-2} \text{tr}(\mathbf{S}_{\underline{\lambda}_q,q}^4) = O(\mathbb{E}V/n)$. Condition (44) follows by an application of the one-sided Chebyshev's inequality for $C_4^2 > \sigma^2 \kappa_q(0, 2) \{1 + o(1)\}$. Since the bound in previous display is $o\{1/\log(n)^2\}$ we further assume that the bound holds uniformly over $q \in Q_n$ with an $o(1)$ term that is $o(1/|Q_n|)$.

8.4.5 Checking condition (45)

Note that if (40) holds then it suffices to check

$$\inf_{f \in \mathcal{W}_\beta(M)} \mathbb{P}_f \left\{ \sup_{\lambda \in L_q} r_n(\lambda, q) \leq K n^{-\beta/(2\beta+1)} \right\} = \inf_{f \in \mathcal{W}_\beta(M)} 1_{\{D_n(f)\}} = 1 + o(1),$$

where $D_n(f) = \{\sup_{\lambda \in L_q(f)} r_n(\lambda, q) \leq K n^{-\beta/(2\beta+1)}\}$.

Proceeding as in Section 8.4.1, the supremum in $D_n(f)$ is attained at $\underline{\lambda}_q$ so if we redefine $R = R_n(\underline{\lambda}_q, q)$ and $r = r_n(\underline{\lambda}_q, q)$, then by the one-sided Chebyshev's inequality $r^2 \leq \mathbb{E}R + \{\alpha/(1 - \alpha)\}^{1/2} (\mathbb{V}R)^{1/2}$. By direct substitution, since $\beta \leq q < \bar{b}$, we have that for sufficiently large n

$$\mathbb{E}R = \frac{\underline{\lambda}_q^{-1/(2q)}}{n} \kappa_q(0, 1) \{1 + o(1)\} \leq K n^{-2\beta/(2\beta+1)}, \quad \text{and} \quad \mathbb{V}R = o\left\{(\mathbb{E}R)^2\right\},$$

for constant K depending on β, \bar{b}, σ^2 , and M . Since the bound on the variance in previous display is $o\{1/\log(n)^2\}$ we further assume that the bound holds uniformly over $q \in Q_n$. Note that in particular, if $q = \beta$,

$$\left[\|f^{(\beta)}\|^2 / \{\sigma^2 \kappa_\beta(0, 2)\} \right]^{1/(2\beta+1)} \leq \left[M^2 / \{\sigma^2 \kappa_\beta(0, 2)\} \right]^{1/(2\beta+1)} = K,$$

uniformly over $f \in \mathcal{W}_\beta$. The statement of the theorem follows.

8.5 Proof of Theorem 4

For simplicity, we assume in the following that $\hat{f}_{\lambda,q}$ is independent of $\hat{\lambda}_f$ (say we make two independent observations at each design point in \mathbf{x} ; this certainly does not impair the performance of the GCV based estimate of λ , and it considerably simplifies the exposition).

Expanding around the oracle λ_f we have, a.s.,

$$\frac{1}{(1 + n\eta_{q,i}\hat{\lambda}_f)^m} = \frac{1}{(1 + n\eta_{q,i}\lambda_f)^m} - m \frac{n\eta_{q,i}\lambda_f}{(1 + n\eta_{q,i}\lambda_f)^{m+1}} O(\hat{\lambda}_f/\lambda_f - 1).$$

Conclude that $\text{tr}(\mathbf{S}_{\hat{\lambda}_f,q}^m) = \text{tr}(\mathbf{S}_{\lambda_f,q}^m) - m \text{tr}\{(\mathbf{I}_n - \mathbf{S}_{\lambda_f,q})\mathbf{S}_{\lambda_f,q}^m\} O(\hat{\lambda}_f/\lambda_f - 1)$, a.s., and that $\mathbf{f}^T \mathbf{S}_{\hat{\lambda}_f,q}^m \mathbf{f} = \mathbf{f}^T \mathbf{S}_{\lambda_f,q}^m \mathbf{f} - m \mathbf{f}^T (\mathbf{I}_n - \mathbf{S}_{\lambda_f,q}) \mathbf{S}_{\lambda_f,q}^m \mathbf{f} O(\hat{\lambda}_f/\lambda_f - 1)$, a.s..

We now determine the asymptotic expectation and variance of risk of $\hat{f}_{\hat{\lambda}_f,q}$. From $n\|\hat{\mathbf{f}}_{\hat{\lambda}_f,q} - \mathbf{f}\|^2 = \mathbf{Y}^T \mathbf{S}_{\hat{\lambda}_f,q}^2 \mathbf{Y} - 2\mathbf{Y}^T \mathbf{S}_{\hat{\lambda}_f,q} \mathbf{f} + \mathbf{f}^T \mathbf{f}$, and the expansions above

$$\begin{aligned} \mathbb{E}[n\|\hat{\mathbf{f}}_{\hat{\lambda}_f,q} - \mathbf{f}\|^2 \mid \hat{\lambda}_f] &= \sigma^2 \text{tr}(\mathbf{S}_{\hat{\lambda}_f,q}^2) + \mathbf{f}^T (\mathbf{I}_n - \mathbf{S}_{\hat{\lambda}_f,q})^2 \mathbf{f} = \\ &= \left\{ \sigma^2 \text{tr}(\mathbf{S}_{\lambda_f,q}^2) + \mathbf{f}^T (\mathbf{I}_n - \mathbf{S}_{\lambda_f,q})^2 \mathbf{f} \right\} \left\{ 1 + O\left(\frac{\hat{\lambda}_f}{\lambda_f} - 1\right) \right\}, \\ \mathbb{V}[n\|\hat{\mathbf{f}}_{\hat{\lambda}_f,q} - \mathbf{f}\|^2 \mid \hat{\lambda}_f] &= 2\sigma^4 \text{tr}(\mathbf{S}_{\hat{\lambda}_f,q}^4) + 4\sigma^2 \mathbf{f}^T (\mathbf{I}_n - \mathbf{S}_{\hat{\lambda}_f,q})^2 \mathbf{S}_{\hat{\lambda}_f,q}^2 \mathbf{f} = \\ &= \left\{ 2\sigma^4 \text{tr}(\mathbf{S}_{\lambda_f,q}^4) + 4\sigma^2 \mathbf{f}^T (\mathbf{I}_n - \mathbf{S}_{\lambda_f,q})^2 \mathbf{S}_{\lambda_f,q}^2 \mathbf{f} \right\} \left\{ 1 + O\left(\frac{\hat{\lambda}_f}{\lambda_f} - 1\right) \right\}, \end{aligned}$$

by Lemma 1, and by Lemma 4. By Lemma 5 we conclude that

$$\begin{aligned} \mathbb{E}\left\{ \mathbb{E}[n\|\hat{\mathbf{f}}_{\hat{\lambda}_f,q} - \mathbf{f}\|^2 \mid \hat{\lambda}_f] \right\} &= \frac{1}{n} \left\{ \sigma^2 \text{tr}(\mathbf{S}_{\lambda_f,q}^2) + \mathbf{f}^T (\mathbf{I}_n - \mathbf{S}_{\lambda_f,q})^2 \mathbf{f} \right\} \{1 + o(1)\}, \\ \mathbb{V}\left\{ \mathbb{E}[n\|\hat{\mathbf{f}}_{\hat{\lambda}_f,q} - \mathbf{f}\|^2 \mid \hat{\lambda}_f] \right\} &= \frac{1}{n^2} o\left\{ \sigma^2 \text{tr}(\mathbf{S}_{\lambda_f,q}^2) + \mathbf{f}^T (\mathbf{I}_n - \mathbf{S}_{\lambda_f,q})^2 \mathbf{f} \right\}^2, \\ \mathbb{E}\left\{ \mathbb{V}[n\|\hat{\mathbf{f}}_{\hat{\lambda}_f,q} - \mathbf{f}\|^2 \mid \hat{\lambda}_f] \right\} &= \frac{1}{n^2} \left\{ 2\sigma^4 \text{tr}(\mathbf{S}_{\lambda_f,q}^4) + 4\sigma^2 \mathbf{f}^T (\mathbf{I}_n - \mathbf{S}_{\lambda_f,q})^2 \mathbf{S}_{\lambda_f,q}^2 \mathbf{f} \right\} \{1 + o(1)\}, \end{aligned}$$

so that again by Lemma 1, and by Lemma 4,

$$\mathbb{V}(\|\hat{\mathbf{f}}_{\hat{\lambda}_f,q} - \mathbf{f}\|^2) = o\left\{ \mathbb{E}(\|\hat{\mathbf{f}}_{\hat{\lambda}_f,q} - \mathbf{f}\|^2) \right\}^2.$$

Now, if $\mathbb{E}(\|\hat{\mathbf{f}}_{\hat{\lambda}_f,q} - \mathbf{f}\|^2) > C\sigma^2 r_n(\lambda_\beta, \beta)^2$, $C > 1$, then by Chebyshev's inequality,

$$\begin{aligned} \mathbb{P}(\mathbf{f} \in \hat{\mathcal{D}}_n) &= \mathbb{P}\left\{ \|\hat{\mathbf{f}}_{\hat{\lambda}_f,q} - \mathbf{f}\|^2 \leq \sigma^2 r_n(\lambda_\beta, \beta)^2 \right\} \leq \\ &\leq \left[1 + \left\{ \mathbb{E}(\|\hat{\mathbf{f}}_{\hat{\lambda}_f,q} - \mathbf{f}\|^2) - \sigma^2 r_n(\lambda_\beta, \beta)^2 \right\}^2 / \mathbb{V}(\|\hat{\mathbf{f}}_{\hat{\lambda}_f,q} - \mathbf{f}\|^2) \right]^{-1} = o(1). \end{aligned}$$

From Sections 8.4.1, $r_n(\lambda, q)^2 = \kappa_q(0, 1)\lambda^{-1/(2q)}/n\{1 + o(1)\}$. By the definition of λ_f (cf. Krivobokova, 2013),

$$\begin{aligned} \mathbb{E}(\|\hat{\mathbf{f}}_{\hat{\lambda}_f,q} - \mathbf{f}\|^2) &\geq \frac{\sigma^2}{n} \text{tr}(\mathbf{S}_{\lambda_f,q}^2) + \frac{1}{n} \mathbf{f}^T (\mathbf{I}_n - \mathbf{S}_{\lambda_f,q})^2 \mathbf{S}_{\lambda_f,q} \mathbf{f} \\ &= \frac{\sigma^2}{n} \left\{ 2 \text{tr}(\mathbf{S}_{\lambda_f,q}^2) - \text{tr}(\mathbf{S}_{\lambda_f,q}^3) \right\} \{1 + o(1)\} = \frac{\sigma^2}{n} \left\{ 2\kappa_q(0, 2) - \kappa_q(0, 3) \right\} \lambda_f^{-1/(2q)} \{1 + o(1)\}. \end{aligned}$$

From (17) and (27),

$$\lambda_{\beta}^{-1/(2\beta)} \geq \{\kappa_q(1, 2)/\kappa_{\beta}(0, 2)\}^{1/(2\beta+1)} \lambda_f^{-1/(2q)} \{1 + o(1)\}.$$

By basic properties of the gamma function, for any $\beta > 1/2$,

$$2\kappa_q(0, 2) - \kappa_q(0, 3) > \kappa_{\beta}(0, 1) \{\kappa_q(1, 2)/\kappa_{\beta}(0, 2)\}^{1/(2\beta+1)}, \quad q \in [\beta/2, \beta],$$

so that the probability that \mathbf{f} belongs to $\hat{\mathcal{D}}_n$ is indeed $o(1)$.

References

- Babenko, A. and E. Belitser (2010). Oracle convergence rate of posterior under projection prior and bayesian model selection. *Mathematical methods of statistics*, 219–245.
- Belitser, E. and F. Enikeeva (2008). Empirical bayesian test of the smoothness. *Mathematical Methods of Statistics* 17(1), 1–18.
- Belitser, E. and S. Ghosal (2003). Adaptive Bayesian inference on the mean of an infinite-dimensional normal distribution. *Ann. Statist.* 31(2), 536–559.
- Belitser, E. and B. Levit (2003). On the empirical bayes approach to adaptive filtering. *Mathematical Methods of Statistics* 12(2), 131–154.
- de Jonge, R. and J. van Zanten (2010). Adaptive nonparametric Bayesian inference using location-scale mixture priors. *Annals of Statistics*, 3300–3320.
- de Jonge, R. and J. H. van Zanten (2012). Adaptive estimation of multivariate functions using conditionally Gaussian tensor-product spline priors. *Electron. J. Stat.* 6, 1984–2001.
- Genovese, C. and L. Wasserman (2008). Adaptive confidence bands. *Ann. Statist.* 36(2), 875–905.
- Ghosal, S., J. Ghosh, and A. van der Vaart (2000). Convergence rates of posterior distributions. *Ann. Statist.* 28, 500–531.
- Ghosal, S., J. Lember, and A. van der Vaart (2008). Nonparametric Bayesian model selection and averaging. *Electron. J. Stat.* 2, 63–89.
- Ghosal, S. and A. van der Vaart (2007). Convergence rates of posterior distributions for non-i.i.d. observations. *Ann. Statist.* 35(1), 192–223.
- Giné, E. and R. Nickl (2010). Confidence bands in density estimation. *Ann. Statist.* 38(2), 1122–1170.
- Johnstone, I. and B. Silverman (2005). Empirical Bayes selection of wavelet thresholds. *Annals of Statistics*, 1700–1752.
- Kimeldorf, G. and G. Wahba (1970). A correspondence between bayesian estimation on stochastic processes and smoothing by splines. *Ann. Math. Statist.* 41, 495–502.
- Knapik, B., A. van der Vaart, and J. van Zanten (2011). Bayesian inverse problems with Gaussian priors. *Annals of Statistics*, 2626–2657.

- Knapik, B. T., B. Szabó, A. van der Vaart, and J. van Zanten (2013). Bayes procedures for adaptive inference in inverse problems for the white noise model. *Probability Theory and Related Fields*, 1–43.
- Kohn, R. and C. F. Ansley (1987). A new algorithm for spline smoothing based on smoothing a stochastic process. *SIAM. J. Sci. Statist. Comp.* 8, 33–48.
- Kotz, S. and S. Nadarajah (2004). *Multivariate t-distributions and their applications*. Cambridge University Press.
- Krivobokova, T. (2013). Smoothing parameter selection in two frameworks for penalized splines. *J. R. Statist. Soc. B* 75, 725–741.
- Low, M. (1997). On nonparametric confidence intervals. *Ann. Statist.* 25(6), 2547–2554.
- Lukas, M. (1998). Assessing regularised solutions. *J. Austral. Math. Soc. Ser. B.* 30, 24–42.
- McAuliffe, J., D. Blei, and M. Jordan (2006). Nonparametric empirical Bayes for the Dirichlet process mixture model. *Statistics and Computing*, 5–14.
- Picard, D. and K. Tribouley (2000). Adaptive confidence interval for pointwise curve estimation. *Ann. Statist.* 28(1), 298–335.
- Robbins, H. (1955). An empirical Bayes approach to Statistics. In *Proc. 3rd Berkeley Symp. on Math. Statist. and Prob.*, Volume 1, Berkeley, pp. 157–164. Univ. of California Press.
- Shen, W. and S. Ghosal (2015). Adaptive bayesian procedures using random series priors. *Scandinavian Journal of Statistics (to appear)*.
- Shen, X. and L. Wasserman (2001). Rates of convergence of posterior distributions. *Annals of Statistics*, 687–714.
- Speckman, P. (1985). Spline smoothing and optimal rates of convergence in nonparametric regression models. *Ann. Statist.* 13, 970–983.
- Speckman, P. and D. Sun (2003). Fully Bayesian spline smoothing and intrinsic autoregressive priors. *Biometrika* 90(2), 289–302.
- Szabó, B., A. van der Vaart, and H. van Zanten (2014). Frequentist coverage of adaptive nonparametric Bayesian credible sets. *Ann. Statist. (to appear)*, 1–42.
- Unser, M. and T. Blu (2000). Fractional splines and wavelets. *SIAM review* 42(1), 43–67.
- van der Vaart, A. and J. van Zanten (2008). Rates of contraction of posterior distributions based on Gaussian process priors. *Ann. Statist.* 36(3), 1435–1463.
- van der Vaart, A. W. and J. H. van Zanten (2009). Adaptive Bayesian estimation using a Gaussian random field with inverse gamma bandwidth. *Ann. Statist.* 37(5B), 2655–2675.
- Wahba, G. (1985). A comparison of GCV and GML for choosing the smoothing parameter in the generalized spline smoothing problem. *The Annals of Statistics*, 1378–1402.
- Wahba, G. (1990). *Spline models for observational data*, Volume 59. SIAM.

- Yue, Y. R., D. Simpson, F. Lindgren, and H. Rue (2014). Bayesian adaptive smoothing splines using stochastic differential equations. *Bayesian Anal.*, 397–424.
- Zhang, C.-H. (2005). General empirical Bayes wavelet methods and exactly adaptive minimax estimation. *Annals of Statistics*, 54–100.